Highly efficient cooling in low-airspeed environments

By Barry Dagan

The physical dimensions of embedded applications are continuing to shrink as users attempt to compress as many functions as possible into a given space. Selecting thermal-management products for these power-hungry miniaturized embedded applications has become a highly complex process. Because embedded applications are placed in a wide array of systems and applications, the thermal-management specification process must be based on a worst-case scenario, both in terms of the magnitude of available airspeeds and the direction from which airflows originate. Designers can choose from a handful of heat-sink options to cool devices efficiently in a low-airspeed environment.

Selecting a heat-sink configuration

When selecting heat sinks for embedded applications placed within low-airspeed environments, designers must consider two important conceptual issues: the density of the pin or fin array and the orientation of the fins.

The first factor that should be addressed is the density of the heat sink. When searching for cooling solutions, engineers often look for heat sinks that possess the largest possible surface area. The common belief is that the more surface area a heat sink has, the more powerful it will be. While this may be true for high-airspeed environments, in low-airspeed environments and the natural convection mode, heat sinks with sparse configurations are generally more efficient.

When densely packed pin arrays are placed in low-airspeed environments, friction forces prevent slow-moving air from breaking the stagnant boundary layers around pins. As a result, surrounding airflows cannot flush through the heat sink, leading to poor convective heat transfer. On the other hand, when sparsely configured heat sinks are placed in low-airspeed environments, the large spacing between the pins presents less resistance to slow-moving airstreams and allows them to penetrate the array of pins, resulting in a more efficient convective heat transfer.

The second factor that should be addressed is fin orientation. Choosing an omnidirectional heat sink is essential for low-airspeed environments, specifically in cases where it is impossible to control the location of the heat sink in relation to the fans in the system. An omnidirectional heat sink ensures the orientation of the air source does not affect performance. Conversely, selecting a unidirectional heat sink introduces significant uncertainty as the orientation of the fan has a substantial impact on performance. For example, in a pin-fin design (see Figure 1), the direction from which the air is approaching does not affect the performance of the heat sink.

Increased demand for highly efficient heat sinks in low-airspeed environments has driven thermal-management vendors to offer new and more efficient heat-sink designs. One of the more innovative heat-sink technologies introduced in recent years is splayed pin-fin heat-sink technology (see Figure 2).

Splayed pin-fin heat sinks contain an array of round pins bent outward in a gradual fashion. The splayed design offers a breakthrough in thermal management as the heat sinks are sparsely configured yet still contain a large surface area. This rare combination stems from the unique splayed design, which takes advantage of generally unused space in the box.

Heat-sink optimization

Cooling power generated in low-airspeed environments is limited, even if the most suitable heat sinks are selected. A common rule of thumb states that a heat sink will be five times more powerful in a high-airspeed environment than if placed in a low-airspeed environment.

Considering the heavy loads of heat being dissipated by today's leading-edge embedded computing systems and the limited space allocated for cooling purposes, the heat-sink specification process has become a challenging task. The following methodologies can help designers improve cooling in low-airspeed environments.

Custom heat sinks

Selecting an off-the-shelf heat sink is always a preferable solution; however, it is not always a feasible one. In cases where

Figure 1

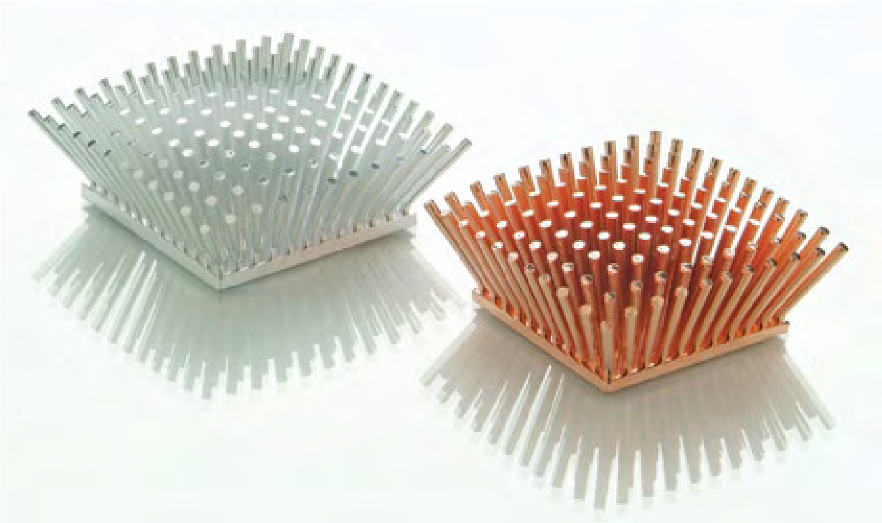

Figure 2

The splayed design offers a breakthrough in thermal management as the heat sinks are sparsely configured yet still contain a large surface area.

standard heat sinks do not provide the desired thermal resistance, a custom heat sink may be required. Selecting a custom heat sink that fills all the available space around the device achieves a significantly more efficient solution because of two factors: the increase in total surface area and the elimination of air bypass.

It is evident that a larger heat sink will offer more surface area, which in turn will increase the cooling power. However, the low pin density of the heat sink should be maintained and the additional surface area should be generated by a larger footprint or from taller pins.

The second less obvious but highly important factor is that maximizing the size of the heat sink can reduce the bypass air typical to low-airspeed environments. Eliminating possible air bypass forces the moving air to flush through the pins, thus improving the air convection throughout the entire heat sink.

Switching to copper

Copper is a significantly more conductive element than aluminum, and as a result, copper heat sinks provide a cooling premium over identically structured aluminum heat sinks. One of the attractive aspects of switching to copper from aluminum is that additional cooling can be derived without more application space.

Copper heat sinks generally outperform identically structured aluminum heat sinks by 5 to 25 percent. The extent of the cooling premium is a factor of several variables. However, the larger the difference between the footprint of the heat sink and the footprint of the device, the larger the performance boost derived by switching to a copper heat sink. This is because copper's highly conductive nature enables it to spread heat along the heat sink's base much faster than aluminum. An additional advantage of using copper in such instances is the prevention of hot spots at the chip's junction.

To demonstrate how switching to copper can provide a substantial cooling premium, consider the process of cooling a powerful, exposed, silicon flip-chip CPU. A typical high-end CPU has a 0.5" x 0.5" exposed silicon heat source and dissipates between 20 to 50 W. Power levels of this magnitude require a relatively large heat sink. Using copper in this case would enable the heat to spread efficiently across the heat sink. In addition to providing approximately 25 percent more cooling power, copper would prevent excessive peak of the junction temperatures.

Fan sinks

As previously discussed, low-airspeed cooling is limited by nature, and in certain situations, cooling hot devices may be impossible without the aid of additional air sources. One of the simplest ways to approach situations in which heat sinks cannot provide the required cooling power is using fan sinks.

Fan sinks are a combination of a heat sink and a fan, with the fan mounted on top of the device in the impingement-cooling mode. Though fan sinks provide substantial cooling premiums, several factors can limit the adoption of a fan-sink solution. First, a current source is not always available. Cost and mean time between failure also can be prohibitive.

To illustrate the performance advantage of a typical fan sink, compare the thermal resistance of a fan sink to a sparsely configured pin-fin heat sink, both placed in a 200 LFM low-airspeed environment. In this example, the overall dimensions of the heat sink are comparable to the overall dimensions of the fan sink. The footprint of the pin-fin heat sink is 1.5" x 1.5" with a height of 0.6", and the fan sink measures 1.5" x 1.5" x 0.3" with a 1.6" x 1.6" x 0.3" fan on top, resulting in an overall height of 0.55". The thermal resistance of the pin-fin heat sink is 2.2 °C/W, while the thermal resistance of the fan sink is 1.35 °C/W, illustrating the clear advantage of adding fan airflow.

Addressing cooling challenges

As the devices used in embedded applications continue to heat up, the need for cooling power will continue to escalate. Engineers designing embedded computing systems will undoubtedly face increasingly more complex cooling challenges, which will lead to the proliferation of copper heat sinks as well as fan sinks. ECD

***Barry Dagan** is technical director of the chief technology office at Cool Innovations, Inc. Barry has been involved in thermal management for more than 25 years in a variety of positions and industries. He holds multiple patents and designed a number of breakthrough thermal-management products. Barry earned his BSC in Mechanical Engineering at the Technion Institute of Technology.*

To learn more, contact Barry at:

Cool Innovations, Inc.
141 Cidermill Avenue, #2
Concord, ON, L4K 4G5
Canada
905-760-1992
barryd@coolinnovations.com
www.coolinnovations.com

efficient cooling

