

[Search this site](#)
or [the web](#)



What's New

Site search
Web search

Tutorial 79 - - December 2007

[PDF Version](#)

[Suppliers](#)
[Tutorials](#)
[Literature](#)
[Books](#)
[News](#)
[Photos](#)
[Newsletter](#)
[Updates](#)
[DrFlipchip](#)
[Home](#)

Bump Cooling

George A. Riley
FlipChips Dot Com

A new cooling approach addresses the growing challenge of heat removal from high performance chips. In October, Nextreme announced integrating thin film thermoelectric coolers (eTEC) into copper pillar flip chip bumps. This combination could eliminate die "hot spots" which often limit IC cooling.

Smaller design features leading to higher-density, higher speed chips are making cooling a major constraint to improving chip performance. As high-performance chips dissipate more power in less space, keeping chip junction temperatures within safe operating limits becomes increasingly difficult. The ITRS semiconductor technology roadmap (1) shows thermal design power requirements for high-performance microprocessors reaching 198 watts in 2008, an 18% increase over 2006.

Intel in a 2005 technical paper identified localized heat removal from chips as a potential cooling approach for their next-generation microprocessors. (2) The maximum on-chip junction temperature must be limited to 90 – 110 °C for performance and reliability assurance. Power consumption, and therefore heat generation, can vary widely at different places on a single chip. The relatively low thermal conductivity of silicon sustains wide temperature variations across the chip. Thus, cooling to a satisfactory average chip temperature is no insurance against unacceptable hot spots. Conversely, cooling the whole chip primarily to reduce hot spot temperature is inefficient. Intel included thermoelectric cooling among their thermal management candidates.

The general approach to thermal management has been back-side cooling of the die. A heat spreader and heat sink attached to the back side of the die transports heat to an air flow or other coolant. This uniformly cools the whole die, without regard for hot spots. Cooling adequate to limit hot spots then would require a larger heat sink. Anticipated future increases in heat generation would require even larger and more complex heat sinks, incompatible with shrinking system size.

In 2006, Nextreme augmented the heat sink approach with embedded thin-film thermoelectric coolers for back-side cooling of packaged die. (3). The thermoelectric cooler works by the [Peltier effect](#), where passing an electric current through a pair of junctions between dissimilar materials such as N-type and P-type silicon creates a temperature difference between the two junctions, so that one junction becomes relatively hotter and the other relatively colder.

Nextreme soldered the hot junction side of the eTEC to the heat spreader, with the cold junction side connected to the back of the die over a hot spot. When current is

passed through the eTEC, it acts as a high thermal conductivity path from the die to the heat spreader, reducing the local temperature at the hot spot.

Nextreme's new approach is front-side local cooling, using the thin-film cooler as the lowest layer of a copper pillar bump. (Fig 1) The cooler removes heat from the low thermal conductivity silicon and delivers it to the high thermal conductivity copper pillar. Separate electrical connections to the cooler provide the drive current. A 60 micron high "thermal pillar bump" is said to sustain a 60 °C temperature difference across its height.



Figure 1. Nextreme thermal bumps. (Nextreme photo).

Front-side cooling faces more limitations than back-side cooling. Backside coolers can be placed anywhere there are hot spots, and sized as needed, without reducing active die area. Front-side coolers must be squeezed into the active die area, among the several thousand bumps of a typical CPU, with restrictions on size and location. (Fig 2) In addition, the die first-level metal must now accommodate the power and ground connections for the coolers. eTEC power adds to the overall package heat budget.




Figure 2. Thermal bumps (large diameters) and normal interconnection bumps on a die surface. (Nextreme photo).

The Nextreme announcement leaves several questions unanswered. Carrying the heat away from the die is step one – but where will the heat go, and how will it get there? Bumps normally attach to a substrate. Will heat be routed to or through a multi-layer substrate or a package bottom? If front-side cooling is merely supplementing backside cooling, the front-side disposal problem may be minimal. However, the effective equivalent of a heat sink will still be required somewhere to transfer the total heat, whether to air flow, fluid microchannels, or my lap.

Nextreme's web site claims that only a layout change is required to add thermal pillar bumps to a chip. Company representatives refused to answer my questions or provide details. The thermal pillar bumps are presently in reliability testing, with pilot production scheduled for the 4th quarter of 2007. Perhaps more information will be available at that time.

REFERENCES

(1) International Technology Roadmap for Semiconductors 2006 Update, Table 6.

(2) Ravi Prasha et.al., "Nano and Micro Technology-Based Next-Generation Package--
Level Cooling Solutions," Intel Technology Journal, Volume 9, Issue 4, 2005

(2) Bob Conner, "Taking the Edge off Hot Spots," *Advanced Packaging*, March 2007.

FOR MORE INFORMATION

Contact [Nextreme](#)

FlipChips Dot Com
210 Park Ave #300
Worcester, MA 01609 USA

Phone 508-753-3572
FAX 508-757-0220
EMAIL DrRiley at flipchips.com

[Site Map](#) | [Contact](#) | [Privacy](#)

© 2007 G. Riley -- All rights reserved