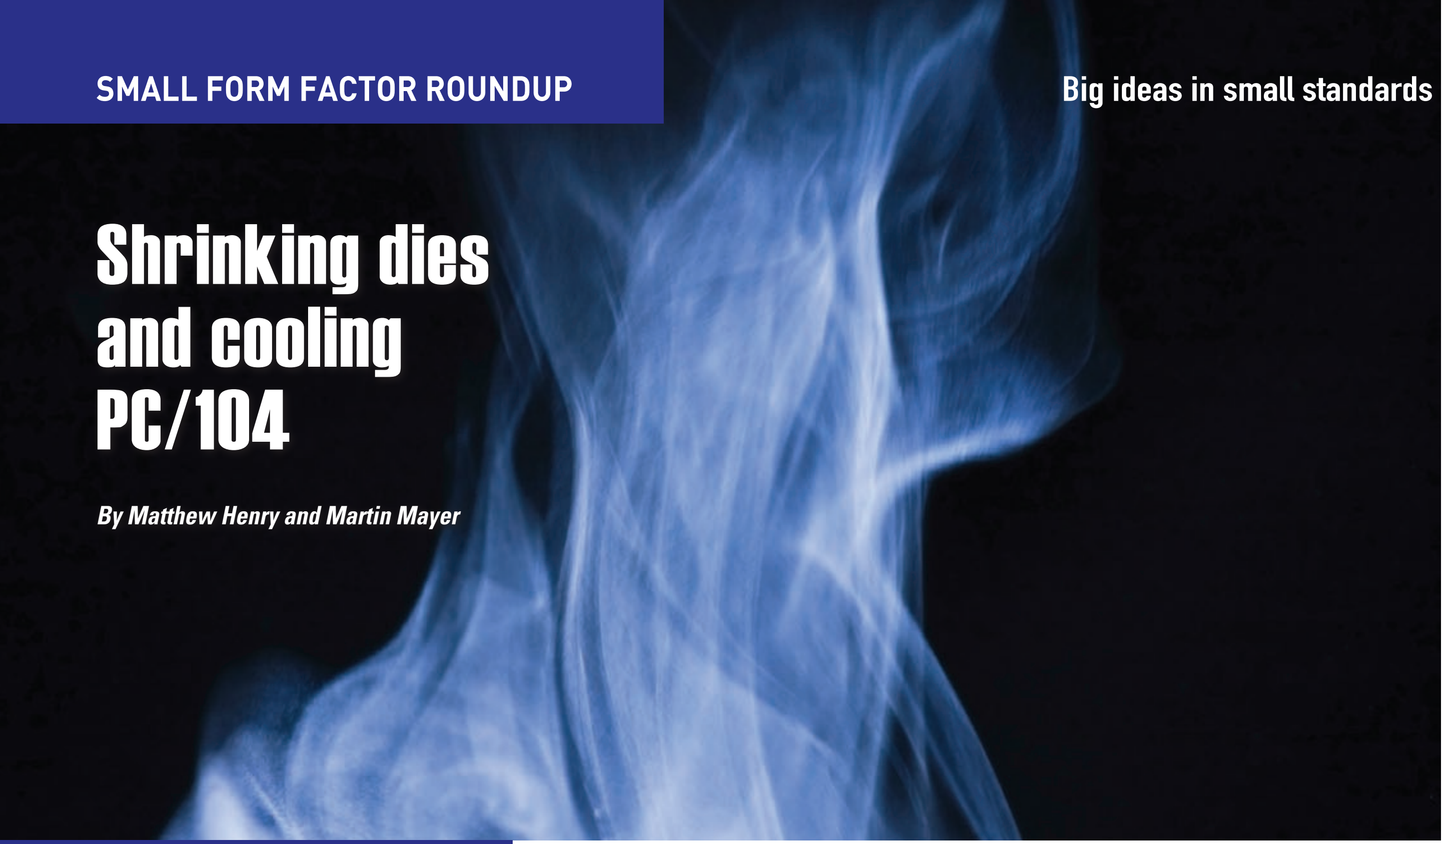

Shrinking dies and cooling PC/104

By Matthew Henry and Martin Mayer

One of the dilemmas facing mechanical engineers is the perpetually shrinking size of processor dies, which presents major thermal management challenges in PC/104 systems. Instead of relying on traditional heat sink and fan solutions, designers can use transmissive thermal interface systems to transfer heat out of small die processor boards in keeping with PC/104 height specifications.

Since the PC/104 form factor's inception in 1992, PC/104 embedded computers have been deployed in less-than-ideal conditions such as low-power enclosed environments and locations where access and serviceability are at a minimum.

Embedded processor technology in the mid-1990s ensured relatively low power consumption and low thermal dissipation, which allowed PC/104 embedded computers to operate virtually maintenance-free using little more than a COTS glue-on heat sink with high-quality radiative properties. Advances in processor technology and an increasing variety of embedded applications demanded new designs that enabled high-performance processors to proliferate the embedded computer market segment. Through the end of the last century, higher performance meant higher power consumption, which translated into higher thermal dissipation and the need for efficient thermal solutions.

In the early part of the last decade, processor die sizes stayed more or less unchanged, while the number of transistors dramatically increased. While the task of cooling the faster, more advanced processors posed a challenge, the die size still provided enough surface contact between the processor and thermal solution; thus, adding a larger COTS thermal solution was sufficient to keep the faster processors cool. As processor manufacturing techniques improved, transistor counts continued to increase while die sizes remained relatively the same. Thermal dissipation then became an issue at the forefront of embedded system deployment. The days of COTS cooling solutions were at an end.

Transmitting heat through the interface

Today, engineers are looking for new and inventive ways to rapidly transfer heat out of small die processors. The available thermal contact area on Intel's Atom Z510 processor is a scant 26 mm^2, roughly the same area as a pencil eraser. The processor dissipates 2 W on average, with a thermal flux of 7.69 W/cm^2. The die's saturation time is a fraction of the saturation time for a processor with a 90 mm^2 thermal contact area that dissipates 5.5 W on average with a thermal flux of 6.11 W/cm^2. Do the math, and you'll see why thermal solutions must be able to wick heat from a processor faster than ever before.

A fan and heat sink solution will go a long way in an open-air environment, but rugged systems aren't always deployed with mechanical air movers and access to available air. While most projects start out on the bench with traditional heat sink and fan combinations, they rarely survive past the proof-of-concept stage and are often retained for development and sustaining engineering functions.

As an alternative to traditional heat sinks and fans, transmissive thermal interface systems allow heat to move from the source (processor or GPU) to a thermal sink, such as an enclosure's chassis wall or an enhanced thermal pathway. Whereas the material used to build a thermal solution or transmissive path does not make a difference, the I/O physical interface is critical. It will not matter that 11000 copper is 166 Watts per meter per degree Kelvin (Wm-1°K-1) better than 1100 aluminum; if the die does not make contact with the thermal interface, the heat will not flow.

Precision is the key to success for thermal contact with a passivated die paddle FCBGA package. As exposed die attach surfaces begin to approach optical flatness, the thermal solution interface with the die paddle must be as flat as realistically possible, regardless of the material chosen.

The example of a regular right cylinder measuring 0.500" length and 0.750" diameter represents a common mechanical input. Roughly 20 to 50 percent of the mechanical input's surface area is used for contact with the CPU die. Contact is dependent on the CPU die area and increases in higher-power, larger-cache, dual-core units. Using 6061-T6 aluminum, 20 percent contact with the input area can handle more than 150 W thermal transmission, making this a suitable initial interface material.

In some cases, the thermal input zone must be enhanced to allow other tolerances to vary without affecting the interface solution's thermal performance. Transport energy at the input can be reduced by replacing aluminum with interference-fit copper. The high pressure associated with an interference fit enhances the transfer of heat from copper to aluminum. While co-milling produces a large, flat surface when necessary, the copper enhancement typically protrudes from the surface and is milled to the required height in a precision CNC milling system. This step more than doubles the thermal transfer capacity, allowing the same energy to be transferred for half of the temperature differential.

While basic accuracies of 0.001" and tolerances of ±0.005" are commonplace in commercially milled products, precision metalwork is required to produce a thermal interface product to tolerances of better than 0.0005". Coupled with precision standoffs, passive thermal interface solutions offer optimal thermal transmission through the interface and provide a precision base for the final PC/104 module stack. Using precision standoffs throughout the stack mitigates additive tolerance creep, preventing stress fractures and delamination of fiberglass boards.

Simulating thermal performance

To help designers get a handle on transmissive cooling solutions and upside-down, CPU-against-the-wall design constraints, Advanced Digital Logic developed the CS-100 Chassis Simulator (Figure 1), which has effective radiative capacity as a 220 in^2 chassis constructed from 0.250" thick K100 or MIC-6 Power Cast Aluminum. The simulator is drilled with the meta-constellation of 3.2 mm recessed mounting holes, allowing for transmissively equipped PC/104 CPUs to be mated to the simulator plate for benchtop operation, thermal stress evaluation, and most importantly, as a thermal reference platform.

Figure 1 | The CS-100 provides a precision thermal interface surface for benchtop chassis simulation.

The chassis simulator represents the finished chassis milling and drill pattern (see Figure 2) and allows easy migration from one CPU class to another. For example, an S15 CPU fitted with a transmissive thermal interface uses the four standard PC/104 mounting holes and one of two different sets of thermal attach screws. Two horizontal M3 holes are provided for cross-compatibility with the 855 product line, and two diagonal 4-40 holes provide compatibility with the LX8 thermal interface, allowing the same chassis hole pattern to be used without modification. Reusing the mounting holes across the entire family of PC/104 products helps minimize risk and improves ROI for designers building products with life cycles that exceed the lifespan of a particular PC/104 CPU.

Some long-life products are upgraded and retrofitted numerous times. The PC/104 platform has demonstrated a longevity that far exceeds the life of any one PC/104 product. Extending that same bolt-pattern philosophy to the thermal attach plate gives designers an upgrade path for their PC/104 stack-based products that is as conceptually simple as upgrading the rims on a car.

Space-constrained applications that use 0.600" PC/104 standard interboard spacing between the chassis wall and the CPU are supported with standard brick-type thermal interfaces. For interconnects that become problematic at standard mating height, customized solutions with critical connectors can switch to right-angle versions, as shown in Figure 1 .

Faster, taller, proof-of-concept brick units with 1.000" board-to-chassis spacing can use Molex SL Series crimp-and-poke connectors (which have a nominal body height of 0.750") on vertical headers without requiring connector modifications. In this type of configuration, several stock CPUs can be evaluated with the CS-100 Reference Chassis, so that the proper CPU speed and feature set can be determined before proceeding down the path of mechanical optimization.

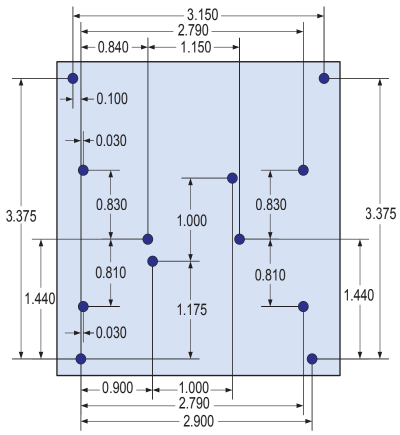

Figure 2 | Standardized chassis mounting holes enable compatibility across all products.

Maintaining the integrity of PC/104

Custom heat solutions vary but hinge on one all-important common element. The maximum available surface contact area must be used to facilitate superior thermal transference from the board to the thermal pathway and out of the enclosure. Keeping close tolerance at thermal junctions is paramount, whether a heat spreader is attached to a processor, chassis, or bulkhead. If this integrity isn't maintained, then the flatness of the contact surface between the processor and thermal transmission medium will be nullified by heat buildup, which will eventually fail the processor when the junctions become saturated.

As die sizes continue to shrink, engineers will continue to design inventive ways of getting the heat out while maintaining the integrity of the PC/104 height specifications. PC/104 stacks often expand downward from the host board to allow maximum opportunity for cooling, but not all designers can support this practice. Therefore, thermal solutions must be kept within PC/104 limits to support those with unique space and design constraints. Processors will keep getting faster and smaller, and mechanical engineers will keep scratching their heads and sharpening their pencils as they await the deployment of processors with diamond-infused die paddles. ➤

Matthew Henry *is product manager and verification engineer at Advanced Digital Logic, based in San Diego, California. He is U.S. Navy educated in Electrical Engineering and has worked in the embedded computing market segment for the past 15 years as an FAE/OEM customer engineer and product manager.*

Martin Mayer *is head of R&D for Advanced Digital Logic. With more than a decade of experience in the embedded sector, starting with writing platform-independent management code for echo cancellers to control code for oil-field applications, he joined Advanced Digital Logic in 1999. Martin earned a BS/CS with Math and Physics minors from Sonoma State University in 1995.*

Advanced Digital Logic
858-490-0597
matt@adl-usa.com
martin@adl-usa.com
www.adl-usa.com