Thermal trends for VXS and VPX deployment

By Bob Sullivan

As serial fabric-based VXS and VPX VME products find their way into the market, component and thermal densities are sure to increase. Moreover, multigigabit signals will also generate large amounts of heat far exceeding today's parallel VME designs. Cooling options range from convection to conduction to liquid, with trade-offs present in each system design.

VXS and VPX are emerging industry standards for embedded computing applications. Each has its own strengths, but both are based on standardized high-performance serial switched fabrics. Initial VXS and VPX products are targeted at performance-driven DSP and multiprocessing applications, having lots of digital signal processors, FPGAs, multicore, and/or Cell processors. These boards also often support PMC and/or XMC mezzanines.

This means that these VXS and VPX boards will typically consume lots of power and will also present system level thermal challenges. Insight into the problem as well as some potential solutions will be discussed.

Problem

In order to support high performance applications, the newer VXS and VPX Standards allow higher power per board than VME64x. While VXS maxes out at a mere 200 W per board, VPX supports over 600W per board. Since these are very small boards (6U x 160mm), even 200W is a lot of power to deal with from a thermal perspective (Figure 1).

Figure 1 shows current and projected power dissipation for a variety of VME board types. The area to the right of the gray area shows projections for board types beyond VXS, while the curve transitions represent the technology shifts from VME64, to 64x, to VPX and so on. The three colored areas under the curves represent a continuum of boards, from a common-type version (middle band), to higher power (+3σ) and to lower power (-30σ) variations. Also, VPX (VITA 46) boards are available in either 12 Vdc or 48 Vdc rail options; liquid cooling will be required at anything over about 250 W.

While VXS and VPX boards are not necessarily hotter than their VME64x predecessors, the trend is definitely in that direction. In our experience interfacing with many board vendors, actual VXS boards typically dissipate around 100 W, with some hotter VXS boards coming in at more than 150 W. VPX is new, but with its higher performance fabric interface and gigabit I/O capability, clocking frequencies and significantly higher data rates mean VPX boards' power dissipation will likely be higher than VXS.

Solution

New IC device package types with increased thermal management are decreasing device junction temperatures. Of the various heat paths from the transistor junctions all the way out to the final system, improved packages are the first design trade-off that yield the largest temperature reduction benefit. Thermal designers can now turn their attention from the device to the module to the board, and finally to the system level.

Depending on application requirements, there are various types of cooling. Typical approaches include forced convection air cooling, conduction cooling, liquid cooling, and spray evaporative cooling. Because of their superior thermal characteristics, liquid cooling and spray evaporative cooling are inevitable in the highest power leading-edge applications. Figure 2 shows the relative heat transfer

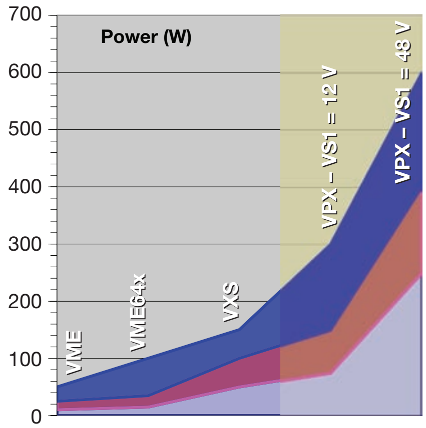

Figure 1

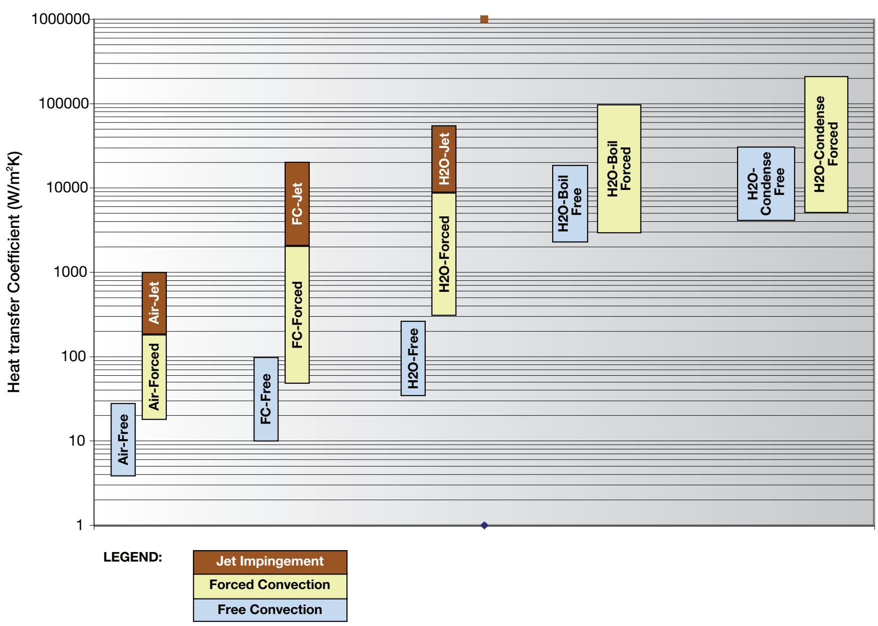

Figure 2

coefficient ranges of the various cooling methods.

Cooling technology advancements are extending the limits for each cooling approach.

Conduction cooling

Hybrid approaches using liquid-cooled side walls can extend the limits of conduction cooling. This approach can cool conduction-cooled VXS or VPX boards at more than 100 W per slot[1].

Technology continues to advance. The emerging VITA 48.2 conduction-cooling standard, which is a companion to VPX, supports improved thermal transfer from the modules to the chassis sidewalls. This, along with low thermal resistance device packages, improved wedge clamps, and novel cold plate designs can extend conduction-cooling performance further.

Forced air convection cooling

Air cooling technology continues to advance in the areas of heat sink design optimization, fan pressure and capacity design, and efficient chassis flow path optimization.

Air cooling at the chassis level is all about mass flow rate (Q) and pressure drop (P). For a given geometry, pressure drop is roughly a function of the square of the mass flow rate. As a result, doubling the mass flow rate results in about four times as much pressure drop across a particular board or chassis topology.

Because of the driving need to increase board functionality, these board designs require larger heat sink areas and result in board level pressure drops increasing as power goes up, compounding the pressure drop and decreasing the mass flow rate. Mezzanine cards increase pressure drop as well. Component density and utilization on the mezzanine card, if not reviewed, will provide an opportunity for creating flow bypass around the devices of concern[2].

The emerging VITA 48.1 REDI air-cooling standard, which is a companion to VPX, will enable better air flow management at the card level, supports a 1" pitch card, which will make it easier to get the heat out through the air flow, and will tend to reduce the pressure drop.

The increase in system power dissipation requires the use of significantly more powerful fans that can handle the higher pressures, as well as reduced pressure drop in the chassis design. These higher performance fans tend to be larger, consume additional electrical input power, and generate more acoustic noise.

Figure 3 shows some actual fan P-Q curves (scaled to a per-slot level) overlaid with some actual board P-Q curves. If there were no other chassis or system-level pressure drops, the operating point would be where the curves cross to give the ideal slot flow rates. As other chassis flow elements are added to provide EMI containment, air filtering, and overall volume minimization, the actual operating crossing points are higher on the fan curve, thus reducing the flow rates.

Liquid cooling

Liquid flow through cooling provides many advantages because it transfers the heat to the fluid very close to the heat generating devices, allowing higher power dissipation for the fluid selected for the system. As platform applications vary, differing fluid types will be evaluated. All fluids used for cooling do not provide the same level of performance when viewed as a group. The Mouromtseff number (see Figure 4) is a measure of the relative heat transfer performance of typical types of fluid used in liquid-cooled applications. These fluid characteristics need to be considered as part of the design architecture strategy.

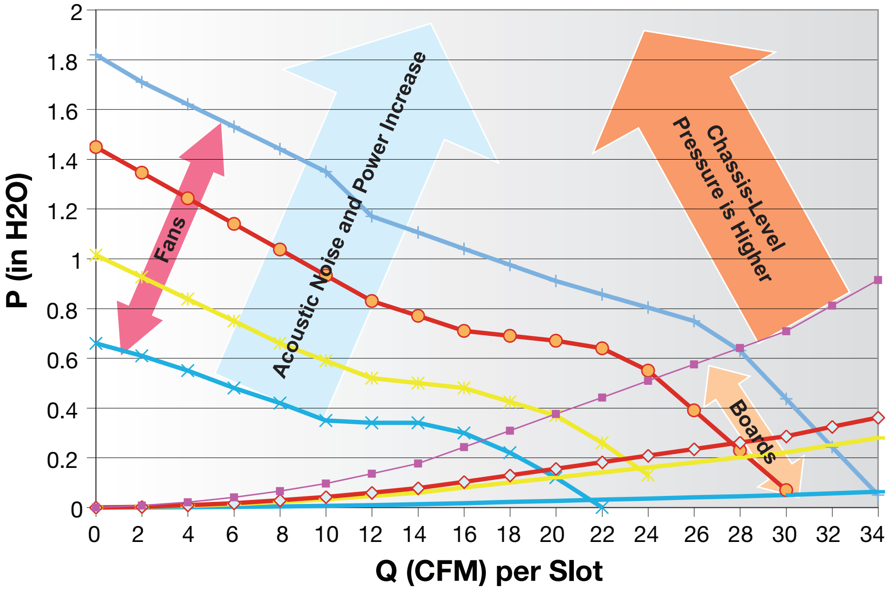

Figure 3

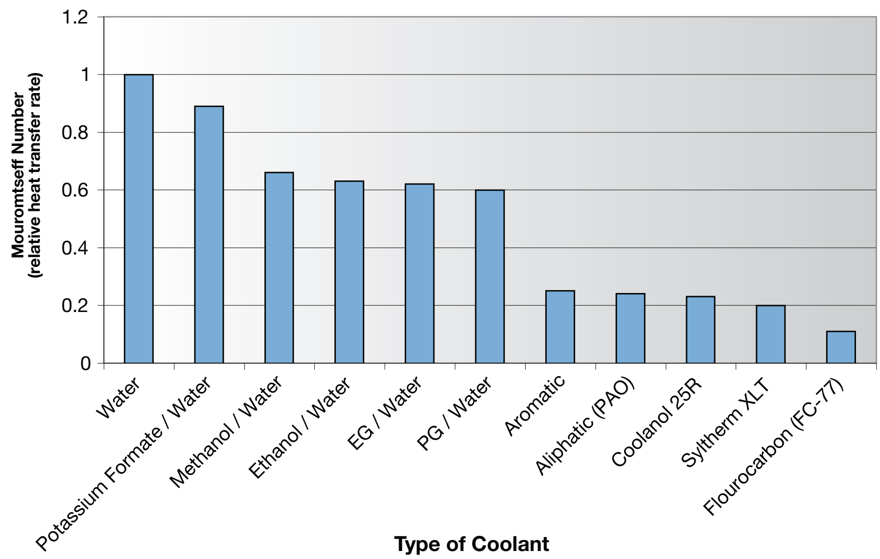

Figure 4

The upcoming VITA 48.3 REDI liquid flow through standard, which is a companion to VPX, adds liquid Quick Disconnect (QD) connectors to the board interface, allowing liquid to flow through each individual board. This is the first industry standard for this type of cooling.

Spray evaporative cooling

Spray evaporative cooling includes a phase change and can provide superior performance. Spray cooling is currently being used in some programs that are typically wet at the chassis level.

VITA 48.3 will enable spray cooling within individual module cold plates. VITA 48.3 is dry at the chassis level; this is a big advantage that will accelerate adoption of spray evaporative cooling approaches.

Technology advancements and cooling

VXS and VPX are stretching the thermal limits of conduction and air cooling. Technology advancements will extend traditional conduction- and air-cooled approaches, but there is a limit to how far this can go. Thermal requirements will inevitably exceed the limits of traditional approaches. Liquid cooling is not the future – it is now! Ω

Robert Sullivan *is Vice President and Corporate Director of Technology at Hybricon, where he has led major programs for communications, military, and industrial applications. He has more than 30 years' experience in the design of high-performance instrumentation and systems and holds a number of patents in the design of high-performance systems. Bob serves on the VITA/VSO and PICMG technical standards committees. He received his MSEE from Northeastern University.*

For more information, contact Bob at:

Hybricon Corp.
12 Willow Road
Ayer, MA 01432
Tel: 978-772-5422
E-mail: bsullivan@hybricon.com
Website: www.Hybricon.com

References:

1. "Deploying High-power Conduction-cooled VXS in Military Applications," *MIL/COTS Digest,* February 2006. Bob Sullivan, Hybricon; Dave Barker, VMETRO; and Jeff O'Reilly, Parker Advanced Cooling Systems.

2. "Using Mezzanine Card Assemblies: Power Dissipation and Air Flow Evaluation," *COTS Journal,* July 2004. Michael R. Palis and Robert C. Sullivan, Hybricon.