

Read About

FPGA cooling challenges

JESD204B

Shrinking technology

Heat density

Wider temperature
requirements

Wider system demands

The Challenge of SWaP for FPGAs

With every new generation of design, we expect something better than the last. The iPhone is a perfect example: with every release, the public is keen to see what new features they can discover. When dealing with rugged, high temperature environments, designing for a better next-generation product has a special set of its own engineering challenges.

Size, weight and power (SWaP) is a cornerstone for many defense applications along with cost and performance. However, these parameters are not always team players, so a choice needs to be made or at the very least compromises for the best overall solution. FPGA-based designs are often seen as the most power hungry designs, and IO devices are getting much smaller while performance is also increasing. As we pack more power into a smaller space, can it be properly cooled? Where are the problem areas? What are the trends? Where should we pay attention? Do you want to buy from a supplier who focusses only functionality?

A prime example of this is high-speed acquisition systems using newer FPGAs: Newer FPGAs offer more resource and higher speed, but does that mean lower power too? Not always. There is often a familiar non-committal phrase: "it depends". FPGA evolution results in lower power per gate between generations due to smaller geometries and new structures such as FinFETs. The power reduction between generations is often cited to be around 30% per gate. These figures come with conditions depending on whether this is static, dynamic or even temperature related. Occasionally power per gate saving is up to 50% power reduction. These power reductions are decent, but newer generation FPGA devices would typically double the gate density too, so if the addition of new algorithms with a similar overall utilization, the board level FPGA design may well end up hotter than its predecessor. In blank terms, this is a pain point.

In this white paper, we'll explore the issue of cooling FPGAs as well as other challenges in FPGA driven high-speed data acquisition systems.




Figure 1: Functionality, higher power and smaller spaces creates system challenges

Info

curtisswrightds.com

Email

ds@curtisswright.com

The Trending Challenges of High-Speed Data Acquisition Systems

Before diving into the realm of cooling issues, there are several other challenges taking trend for FPGA driven high-speed data acquisition systems.

Applications require a certain amount of FPGA resource such as DSP blocks, gate, onchip memory, and flip-flops to implement the algorithms. For some applications, this FPGA resource is the limiting factor, despite having enough external interface resources to physically hook up more ADCs and DACs. For these applications, it is all about the on-chip FPGA resource per channel of IO. However, if the FPGA resource is doubled, there is the option for doubling the IO provided there are enough IO pins to support the IO chips. From generation to generation of FPGA, power needed for IO does not drop by 30% or 50%, as does the average power for on-chip resource. Consequently, the overall power dissipated by a new generation of FPGA can actually increase; This is not taking into account higher inrush currents that need to be handled by the power supply at power on. Although the overall system power might be less for applications that double the IO to go with doubling the FPGA resource (assuming 50% less power per FPGA gate generation on generation), the FPGA would still draw more power. This means cooling for high temperature environments becomes even more challenging using conventional techniques, and this is not a “free” upgrade.

In some scenarios, it may be more effective to spread out the IO and FPGA resource across multiple cards. However, that could come with a cost both for additional cards and larger systems. For a system upgrade, the cost may lie in the rest of the system, so the objective may be to replace the distributed cards rather than to consolidate.

Shrinking Technology

New generations of high-speed ADCs and DACs are adopting high-speed serial interfaces such as JESD204B. Although JESD204B has been around for many years, it has now become the default interface of choice for IO chip vendors because the fastest devices, especially multi-channel devices, would otherwise need an unreasonable number of IO pins. While FPGAs have continued to gain speed, the demands of a similar speed increase in parallel IO pins has not been possible or practical. This has meant

more IO pins on newer FPGAs have been given over to high-speed serial interfaces (and more grounds) across device package sizes. In the past, the number of parallel IO FPGA pins may have been another limiting factor as to whether addition IO channels could have been supported. That is, assuming there was plenty of FPGA onchip resource.

Having so many high-speed serial IO channels on medium and large FPGAs provides the option of having a higher quantity of ADCs and DACs. Today's JESD204B-based high-speed ADCs and DACs may only be a quarter of the package size compared to their parallel interface predecessors due to them only needing a fraction of the IO pins, which in turn defines the package size. A smaller package size in turn means more options for either shrinking the overall board size or putting more onto a given area. However, small is not always easy. A JESD204B part may be a quarter of the device size compared to a parallel version. If the power has not gone down by a similar proportion, then the heat density will increase. In reality the increase in heat density, perhaps by a factor of two, isn't likely to be a great problem, but it is yet another aspect that needs attention in a high temperature environment.

Outlined above are a number of issues associated with newer generations of FPGA driven high-speed IO: higher power dissipation for newer devices, higher heat density for both FPGAs, ADC and DACs, and more IO devices on a newer generation of IO card. This directly relates to more attention needed for cooling.

Heat Density

For a high-temperature environment, 85°C or above is typical for a conduction-cooled cold wall interface and around 70°C for air-cooled. This is manageable for a 20-25W XMC, 90-100W 3U VPX and perhaps 150W 6U VPX format card. At these power levels on these formats, manageable does not mean trivial, but requires careful thermal analysis, design and testing. While it is possible to provide power to FPGAs on these formats, even moderately sized FPGA designs could easily exceed the cooling limits using standard heat frames. This begs the question: If a large FPGA can be fitted, can it be cooled to my target environment? Within these parameters, XMC and VPX are typical card formats used in rugged environments.

Boards powered beyond the levels outlined can be cooled, but require something “more exotic” such as alternative heat frame with much lower thermal impedance or liquid

cooling, which can also include spray cooling. High power cooling accordingly comes with a cost penalty. The VPX community, through the VITA organization, has defined standards based on Air-Flow Through (AFT), otherwise known as VITA 48.8, built upon the VPX format. These cards provide a heat frame with clamps over a VPX card and a slot that permits high quantities of air to be blown across the heat frame within a conduction-cooled chassis. The air does not directly cool the electronics, but the AFT heat frame in a hybrid conduction-cooled or air-cooled setup. This allows a 3U VPX card a power budget 50% or 100% higher, thereby enabling high FPGA resource usage along with high-density IO solutions. AFT does have a cost premium at the system level, but fewer cards saves cost. Costs can also be reduced with a smaller chassis requirement along with the user being able to use more of the available resource for SWaP optimization.

Wider Temperature Requirements

Small form-factor design are often destined for the smaller defense platforms from UAVs to fighter aircraft. For these platforms with low temperature operations, SWaP is even more important. These platforms can be deployed anywhere in the world, from hot deserts to the arctic, and that's just where they may take off and be turned on. Typical low temperature rugged cards have a -40°C operation specification and -55°C storage. Since this storage temperature requirement is below its operational specification, additional operation aspects need to be employed involving either waiting for the equipment to passively warm up or to add in heaters to allow the equipment to turn on after a short wait.

Alternatively, a card or system can self-heat. This can be achieved by designing the equipment to be fully operational down to the extended low temperature. This can be expensive and key devices such as the FPGA may not be available in military or automotive temperature grades. A second alternative is to design a sub-set of the card or system to power up at the lower temperature, self-heat, monitor the temperature through the card and automatically boot the rest of the design.

In the case of FPGA cards, turning on FPGAs under the control of a robust low temperature power supply and controllers is an effective way of self-heating. If the FPGAs are not military grade, they can still be turned on, but cannot be guaranteed to work in any way. In this case, when the temperature has risen to the devices specification, the FPGA can be "rebooted". Careful attention does have to be made during the design due to subtle potential issues such as meta-stability lock-ups. Newer generation rugged design for the defense industry are starting to use designs that can self-heat. The benefit here is system heaters are required and an extended operation temperature range can be supported, providing SWaP optimization to match its target environment.

Export controls can apply to product designed for use below -54°C. Thus, it is possible to design cards for the defense community that are not export controlled, thereby simplifying the sales and supply process. With that said, high-speed ADC and DACs and large FPGAs may be subject to a different set of export controls.

Wider System Demands – Not just a board level view

Newer generation FPGA based designs offer very real benefits. Nevertheless, it is not about what functionality a new design offers, but whether it is usable. Increasingly, more design competence is needed to design products that match the desired environmental requirements. For high performance solution, the power and heat density is going up, not down. Increased development cycles are required to design and thoroughly test large FPGA based solutions and to use what once considered exotic techniques that are destined to become standard build formats. Doing this requires infrastructure, experience and a systems based approach.

Having all parts of a system designed for a rugged high temperature environment instead of only particular components can make or break your program's success. Curtiss-Wright Defense Solutions has particular skills in this environment and a track record in developing, testing and integrating high performance, high power FPGA, DSP, GPGPU and processing solutions. For more information, please browse the links on the next page.

Author



Jeremy Banks, BSc (Hons) in
Electronic and Electrical Engineering
Product Marketing Manager
Curtiss-Wright Defense Solutions

Learn More

Product: [VPX3-530 - 3U VPX Virtex-7 FPGA ADC/DAC](#)

Video: [Meeting the demands of signal intelligence and electronic warfare with the VPX3-530](#)

Case Study: [The Bigger Picture for Radar Sub-systems](#)

White Paper: [Dissecting Rugged Development of COTS FPGA Hardware](#)