

# A 2.5 TFLOPS ATCA Radar Signal Processor

By Richard Jaenicke, Director Market Development, Mercury Systems



## The Design Team

Mercury Systems, Inc. is a domestic US designer of secure processing subsystems serving the defense community

## Challenge

Building a heterogeneous compute blade with -server-class CPUs and a high-performance FPGA with great floating-point performance in single a slot that can be deployed in military ships and wide-body aircraft posed a significant challenge for the design team at Mercury Systems. Additional goals included multi-server system management, reduced cabling, and being based on open architecture for both the hardware and software.

## Solution

The design team engineered two server-class Intel® Xeon® CPUs and a powerful FPGA onto a single ATCA blade, with necessary connectivity, power, and cooling.

## The Project

Airborne and shipborne radar signal processing subsystems are often implemented with a combination of multicore, server-class CPUs and hardware accelerators. This combination can provide multiple TFLOPS of floating-point computational performance using versions of commercially available chips. But these environments impose requirements on volume, power, and ruggedness that preclude commercial solutions. Additionally, military systems are increasingly demanding a level of system management that is unfamiliar in most commercial markets. We set about to bring commercial server-class CPU technology and hardware acceleration—in the form of advanced FPGAs—into the world of military systems.

## The Design Challenge

In the commercial world, accelerated servers are often built using an Intel Xeon E5 motherboard and an FPGA add-in card in a 1U or 2U rackmount chassis. Radar signal-processing systems could exploit this level of performance, but they require significant processing power densely packaged for a harsh environment. The density and ruggedness requirements generally necessitate compromises in computational performance due to the size, power consumption, and package type of the highest performance processors. For instance, smaller, lower power, BGA CPUs such as Intel Core™ i7 client CPUs or mid-range embedded Xeon-D processors are often substituted for full server-class Xeon E5 processors that are larger, higher power, and in an LGA package. Likewise, the FPGA selected for the front-end processing might be restricted in terms of floating point performance and memory size.

## The Design Solution

To reduce the need for such tradeoffs, we chose an AdvancedTCA (ATCA) blade as the form factor. At 322 mm high (8U) x 280 mm deep and 1.2 inches wide, the blade provides enough surface area to fit a large FPGA (45x45 mm) in addition to two 12-core Xeon E5 processors and their supporting chipset while still being able to be ruggedized. To meet the vibration requirements, the Xeon E5 processors are converted from land grid array (LGA) to ball grid array (BGA) and then soldered directly to the blade. An Arria® 10 GX1150 FPGA provides sensor protocol bridging and preprocessing as well as co-processing for the Xeon CPUs.

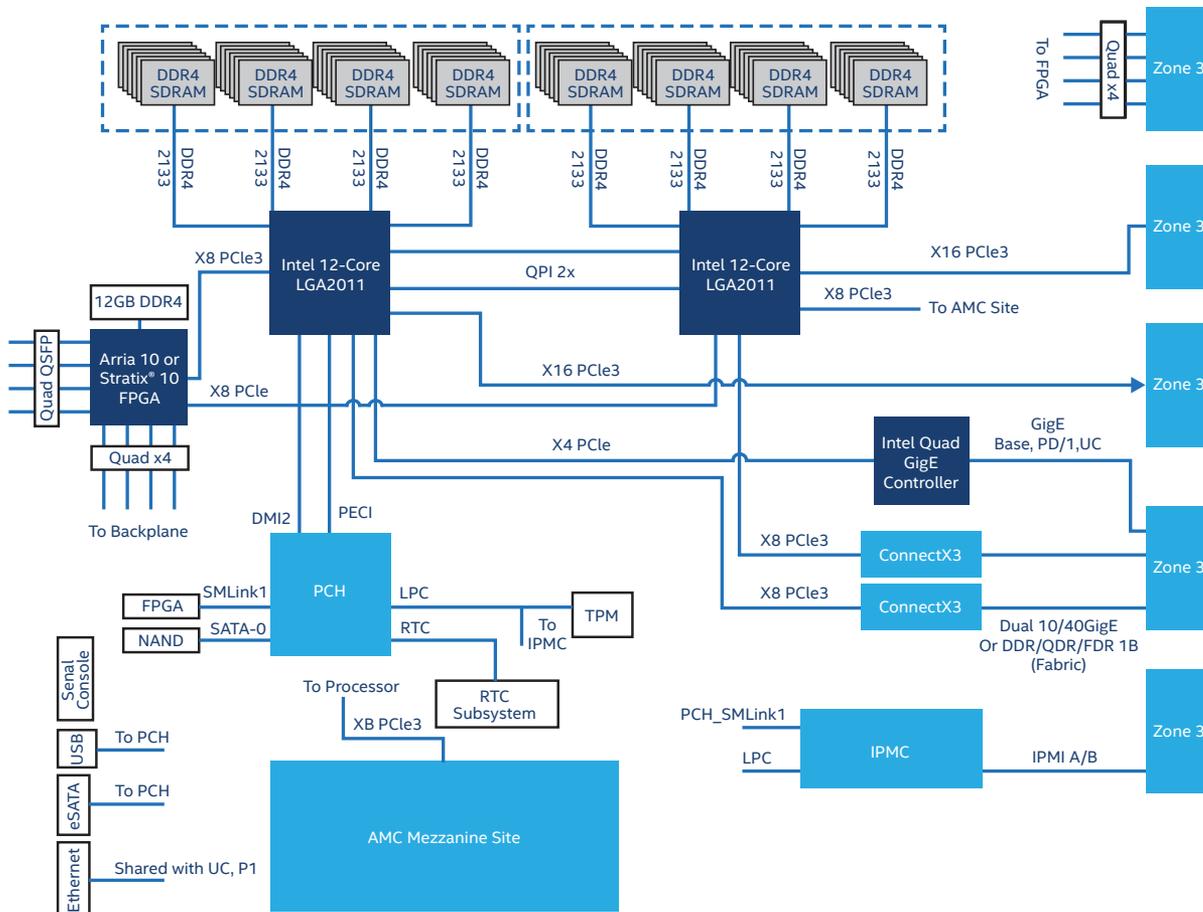
AdvancedTCA blades fit into a “shelf” that typically can hold up to 14 blades in 19 inch rackmount chassis. The system includes a backplane for the fabric interconnect, optional switch blades, and a shelf manager. The shelf manager communicates with the blades over IPMI, and can handle or report error conditions. The backplane connects all the dual processor servers, eliminating piles of complicated cabling. One or two slots in the shelf can be used for switch cards for high cross-sectional bandwidth applications

## Theory of Operation

The sensor signal processing is divided between the Intel Xeon E5 v3 CPUs and the Arria 10 FPGA. In general, the FPGA is used as a preprocessor, but is also available as an offload engine. In the case of radar signal processing, the FPGA does the beamforming and often the pulse compression while the CPUs perform the back-end processing including Doppler processing and target detection.

FPGAs are ideally suited for beamforming, which demands high throughput and low latency while being conducive to a highly parallel implementation. A radar antenna can have thousands of elements tracking hundreds of targets, each requiring processing in two or three dimensions, depending on the type of radar.

Data from radar or electro-optical/infrared sensors enters the Arria 10 FPGA through four Quad Small Form-factor Pluggable (QSFP) transceivers on the front panel of the blade or through four x4 I/O lanes from the backplane connector. The sensor data is usually 40 Gbit Ethernet, but could be a lower-level sensor protocol such as ANSI/VITA 17.1 Serial Front Panel Data Port (SFPDP). Using an FPGA provides



the flexibility to adapt the protocol to legacy sensor outputs. The first task performed by the FPGA is to extract the sensor data from the incoming data packets. The four input channels can be treated as a group or as independent channels. Processing for independent channels can be isolated further if desired to enable separate processing at different levels of sensitivity and security. The algorithms executed in the FPGA typically include tasks such as signal conditioning, filtering, and beamforming. The FPGA has 12GB DDR4 memory available for buffering and intermediate results. The final task of the FPGA is to package up the preprocessed data and send it to each of the two Xeon E5 V3 CPUs over independent x8 PCIe Gen3 links.

The two 12-core processors are linked via two instances with the high-speed, low-latency Quick Path Interconnect (QPI) interface, each of which provides a 38 GB/s (bi-directional) data transfer rate. This interconnected processor architecture is optimized for the intense data movement needed by high performance signal processing algorithms, such as all-to-all corner turn operations where large data sets are written in columns but read in rows. From a software perspective, this QPI architecture allows the back-end processors to be configured with a single kernel NUMA-aware operating system running across both CPUs.

Multiple ATCA blades are connected via a fabric interface over the backplane through two Mellanox® ConnectX-3 host adaptors. Bridging between the native Gen3 PCIe interfaces on the Intel Xeon processors and the AdvancedTCA fabric channel, the ConnectX-3 can be configured to support InfiniBand (DDR, QDR or FDR-10) or 10/40 Gigabit Ethernet as the data protocol.

## Results

The main goal of the solution was to pack as much processing and high-bandwidth I/O into a small, rugged package while keeping the processors reliably cool. The two 12-core Xeon CPUs have a total of 1.38 TFLOPS peak and the Arria 10 GX 1150 FPGA adds another 1.36 TFLOPS. A full-size 14-slot ATCA chassis is 12U high and 23 inches deep and fits into a 19-inch rack. This is 2U shorter than 14 1U servers, but more importantly it saves almost 5 inches in depth. The result is 30% smaller volume while meeting the much higher requirements for ruggedness and operating temperature.

To enhance the ruggedness and reduce the amount of air required to cool the processors when operating at altitude, the blade uses the OpenVPX ANSI/VITA 48.7 Air Flow-By® cooling technique adapted for the AdvancedTCA form factor. Air Flow-By systems cover both sides of the blade with a conformal heatsink that includes fins extending to the remaining permissible volume for the blade. An additional effect of the more efficient cooling is a five-fold improvement in mean-time between failures (MTBF). The encompassing module packaging also provides physical protection, improved ruggedness, and is an effective Faraday cage for higher EMI protection.

