Xeon-D and 3U VPX combine for cognitive EW

MARC COUTURE CURTISS-WRIGHT

The role of FPGAs in cognitive EW

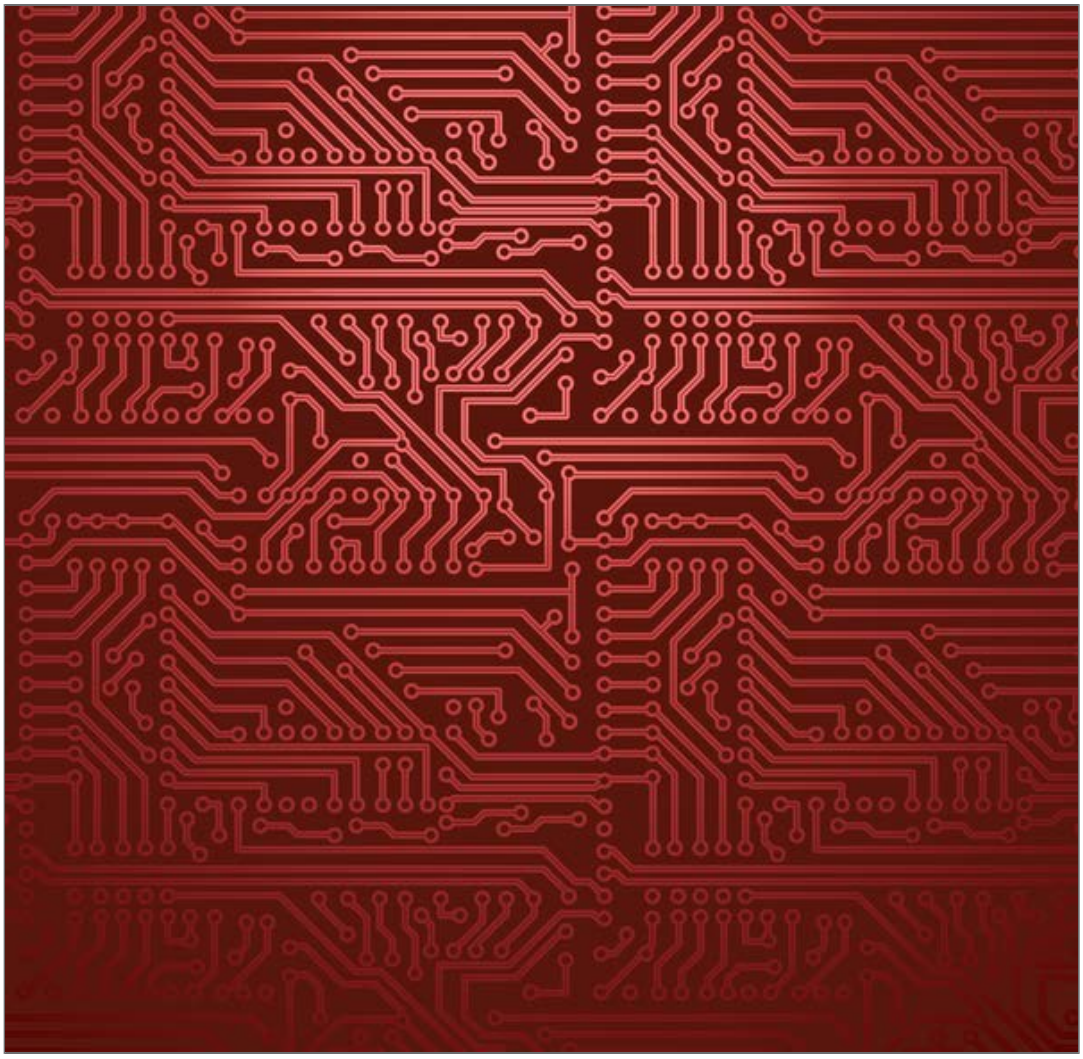

Intel's new Xeon D system-on-chip (SoC)is making large numbers of x86 processing cores readily available for embedded defense applications. With an architecture designed to support math-intensive processing and very-high-bandwidth data transfers, the Xeon D enables advanced cognitive electronic warfare (EW) applications to

operate in small size, weight, and power (SWaP)-constrained platforms.

Unlike conventional radar systems, new software-defined digitally programmable radars are able to generate previously unencountered waveforms that do not match known waveforms and pulse trains already on an EW system's pulse descriptor word (PDW) list. The PDW typically contains all the collected data for a specific pulse, including time of arrival (TOA); angle; pulse width, power, and frequency (superhet); or frequency band. In order to defeat never-before-seen waveforms, system designers are developing a new generation of cognitive EW systems that are able to quickly adapt to changes in the radio frequency (RF) environment and almost instantly make decisions about how to respond to unfamiliar threats.

These cognitive EW systems are well served by the recent introduction of new multicore Intel Xeon processor D (Xeon D) devices that deliver the greater thermal range performance required by SWaP-constrained EW pod environments. Cognitive EW systems built with these new devices promise to provide an alternative to today's EW systems, which are typically implemented with field-programmable gate array (FPGA)-based system-level architectures. While FPGA-centric EW systems that implement digital radio frequency memory (DRFM) and other EW techniques in firmware via VHDL and Verilog are sufficient for intercepting and prosecuting known waveforms and pulse trains on the PDW list, they lack the dynamic flexibility needed to counter new waveforms generated on the fly by sophisticated adversaries using digitally programmable radar.

The role of FPGAs in cognitive EW

That'snot to say that FPGAs don't play an important role in these new cognitive EW systems. Unlike general-purpose processors, FPGAs are very good at sophisticated highly parallelized, high-throughput DRFM

techniques such as range gate pull-off. FPGAs, however, are not very good at making decisions and dynamically changing their own architecture, which is key to cognitive EW. One approach for implementing deployable embedded cognitive EW is to couple a cluster of general-purpose Intel x86 cores – such as the eight, 12, or 16 cores available on a Xeon D – with large FPGAs. In this type of heterogeneous system architecture, the FPGA provides the high-speed parallelization while the Xeon D provides the required real-time supercomputer-class analytic and metadata processing. The decision-making speed of the system's Xeon D cores and the processing performance of the FPGA work in concert to enable the system to respond to unknown waveforms by synthesizing a mix of responses to quickly create the best defense against the new threat. This response may include the partial or complete reconfiguration of the FPGA, depending on the findings of the Xeon D. The extremely low-latency decision-making capability of the x86 cores actually enables the cognitive EW system to select the proper set or mix of signal-manipulation techniques to adapt to the threat while in-theater.

The Xeon D expansion

During the last five to ten years, while FPGAs have dominated EW system development, Intel processors were limited to a maximum of four cores. Previously, x86-based general-purpose processors have served in EW systems, but usually only to provide system management and to handle the man-machine interface. The expanded multicore architecture of the Xeon D enables the devices to serve as active participants in the prosecution of RF emitter stacks. In addition, the recent introduction of these processors in ball grid array (BGA) packages gives EW system designers access to a rugged processor that can provide four, 12, or 16 cores in the same power footprint as the earlier four-core devices. One advantage of the BGA packaging is that the device's entire bottom surface can be used for interconnection pins, which is necessary for high-bandwidth operations. Moreover, each of the Xeon D cores is supplemented by a powerful AVX2 SIMD engine to deliver enough processing power to

execute complex decision-making and high-bandwidth DSP math (Figure 1).

Figure 1: The Intel Xeon processor D-1500 product family features 8-/12-/16-core versions with enhanced performance at low power, suitable for use on rugged open-architecture modules designed for deployment in extremely compute-intensive EW and command, control, communications, computers, intelligence, surveillance,and reconnaissance (C4ISR) aerospace and defense applications in harsh environments.

(Click graphic to zoom)

With the large number of cores at their disposal, designers can flexibly partition and dedicate sets of cores to different EW techniques and dynamically allocate processor resources. The Xeon D, with its multiple closely-linked x86 cores and math engines, is designed to support math-intensive processing and very high bandwidth data transfers and can truly be considered a “mobile server.”

Although there are other processing devices that are capable of delivering higher FLOPS performance, Xeon D delivers an average of 0.5 TFLOPS, depending on the specific chip SKU, in a manageable 35

W to 45 W package. In comparison, power requirements for larger Xeons and GPUs range from 75 W to 100 W and higher, making them extremely difficult to cool even with VITA 48 REDI techniques. Since the Xeon D is a compact SoC, commercial off-the shelf (COTS) vendors can create SWaP-optimized VPX modules that provide enough space to include a 25 W XMC mezzanine. To complete the system, wideband analog-to-digital converters (ADCs) and digital-to-analog converters (DACs), used as microwave IF receivers or even RF tuners and for the digitization of interferometry for jamming, can be integrated onto an XMC mezzanine module that fits above the Xeon D in a one-inch pitch. This setup means that the FPGA and digitizers can be placed in extremely close proximity to the multiple x86 cores on a 3U module, which can improve overall performance and significantly lower latency.

3U VPX for SWaP-optimized cognitive EW

The 3U VPX form factor provides a useful way to deploy new Xeon D/FPGA-based card sets into next-generation cognitive EW pods. What makes the device an especially good fit for SWaP-constrained applications is the fact that, unlike larger and hotter commercial Xeon devices and Core i7s processors, Xeon D eliminates the need for a separate “Southbridge” platform controller hub. Xeon D is actually a single SoC. The serial ATA (SATA) interface and general-purpose I/O functionality used in laptops and servers has been subsumed into the device, making it suitable for use on small-form-factor conduction-cooled 3U modules where real estate is limited. In rugged deployed applications, larger land grid array (LGA)-packaged Xeon processors require special – even exotic – cooling techniques that may not be appropriate for many programs since they add risk and undesirable costs. In contrast, the more affordable Xeon D, with its greater operating range, is able to use standard cooling envelopes and techniques

A new 3U VPX Xeon D solution

An example of a Xeon D-based 3U VPX module is the Curtiss-Wright

CHAMP-XD1 card, a DSP engine module for use in very compute-intensive High Performance Embedded Computing (HPEC) applications (Figure 2). This open-architecture module features high-speed DDR4 memory and bandwidth along the OpenVPX data, expansion, and control planes. It has XMC card expansion and the choice of one Gigabit or 10 Gigabit Ethernet (GbE) interfaces along the control plane.

Figure 2: Curtiss-Wright's 3U OpenVPX CHAMP-XD1 is a SWaP-optimized DSP engine module that leverages the performance specs of Intel's Xeon processor D architecture for EW applications.

(Click graphic to zoom by 1.9x)

Marc Couture is the Senior Product Manager for Intel, PowerPC, and GPGPU based digital signal processors in the ISR Solutions group at Curtiss-Wright. He has worked in the embedded COTS industry for over two decades, having specialized in High Performance Embedded Computing and RF/microwave technologies. Marc is a graduate of the University of Massachusetts Dartmouth with a Master of Science in Electrical Engineering. Readers may reach him at mcouture@curtisswright.com.

Curtiss-Wright www.curtisswrightds.com

THIS ARTICLE WAS PUBLISHED ON FEBRUARY 2nd, 2016.