

MICROPROCESSOR *report*

Insightful Analysis of Processor Technology

LATTICE FIRST WITH FPGA IN FD-SOI

Debut Chip in Nexus Line, CrossLink-NX Aims for Embedded Vision

By Bob Wheeler (February 3, 2020)

Aiming for markets underserved by Intel and Xilinx, Lattice is first to bring the advantages of fully depleted silicon-on-insulator (FD-SOI) technology to FPGAs. The debut product in the new Nexus platform is CrossLink-NX, which Samsung manufactures in its 28nm FD-SOI process. It's designed for embedded-vision systems like the example in Figure 1. Unlike shipping 40nm CrossLink chips, the NX generation includes DSP blocks to accelerate AI inferencing, as well as a PCI Express interface. Both generations have MIPI interfaces for connecting image sensors and SoCs. The expanded capabilities of CrossLink-NX, however, take it from a simple bridge to a video coprocessor.

At the Power/Performance Nexus

Compared with prior Lattice FPGAs, Nexus reduces both dynamic and static power thanks to its use of FD-SOI. Samsung licensed the 28nm FD-SOI technology from ST Microelectronics, which deployed it in 2014 (see [MPR 8/15/16](#), "FD-SOI Offers Alternative to FinFET"). The Nexus design uses a reverse (or back) body bias to reduce leakage current by up to 75% compared with similar products built in 28nm bulk CMOS. Lattice doesn't employ forward bias to reduce threshold voltage; the design employs a conservative 1.0V core voltage.

Another FD-SOI benefit is a large reduction in soft errors caused by radiation, as the channel area is a smaller target for neutrons and alpha particles. Because conventional FPGAs, including Nexus, configure their fabrics using SRAM, a single-event upset (SEU) can cause unrecoverable failures. For mission-critical designs, FPGA vendors provide soft-error-mitigation cores to detect SEUs, but these errors still interrupt real-time operations. Nexus cuts the error rate by about two orders of magnitude, virtually eliminating these interruptions.

The Nexus programmable logic is based on the four-input-LUT (LUT4) design in the 40nm ECP5 generation (see [MPR 5/12/14](#), "Lattice ECP5 Complements SoCs"). To boost performance in AI applications, Lattice increased the new platform's memory-to-logic ratio by adding 512Kb large-memory (LRAM) blocks. To reduce I/O-configuration time at boot, Nexus divides its configuration image (or bit stream) into an I/O section followed by a logic section. It also optimizes the quad-SPI flash-memory interface to load faster by polling for ready and by operating at 150MHz. The company touts the "instant on" CrossLink-NX configuration, which takes 3ms to configure the I/Os and 14ms to configure the entire device.

Embedded Vision Sees 2020

Lattice prioritized CrossLink-NX as the first Nexus design based on customer feedback and on the strategic importance of embedded vision. The new product builds on the shipping CrossLink chips, which are small FPGAs for bridging video and display interfaces. This task usually requires

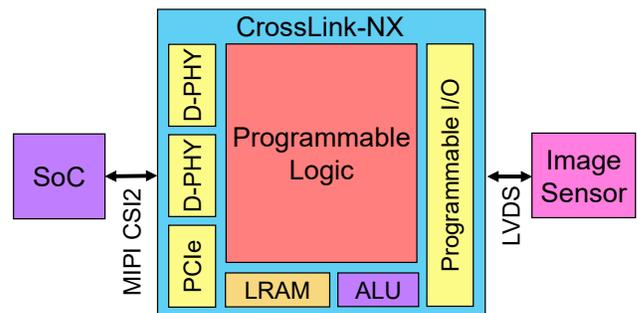


Figure 1. CrossLink-NX edge-AI system. The FPGA can connect multiple image sensors to an SoC while offloading some processing tasks.

connecting a proprietary LVDS camera interface to a MIPI-standard port on a processor or SoC. CrossLink handles MIPI D-PHY, a serial interface with up to four lanes that share a forwarded clock. The shipping products include two D-PHY ports with four lanes running at up to 1.5Gbps per lane. MIPI's Camera Serial Interface 2 (CSI2) protocol operates over D-PHY as well as alternative physical layers.

As Table 1 shows, CrossLink-NX also provides a pair of hardened four-lane D-PHY v1.2 ports, but it achieves up to 2.5Gbps per lane and 10Gbps per port. New to the NX generation is a hardened PCI Express Gen2 x1 interface for connecting SoCs that lack MIPI, particularly for industrial and virtual-reality applications. The first CrossLink-NX device has up to 74 high-performance (HP) programmable I/Os to handle other high-speed interfaces including LVDS, SGMII, and DDR3 SDRAM, as well as up to 12 additional 1.5Gbps D-PHY lanes.

Lattice announced two CrossLink-NX models: the NX-40 has 39,000 logic cells and is sampling now, whereas the NX-17 has 17,000 logic cells and is due in 4Q20. The latter has 2.5x as much LRAM to compensate for its smaller embedded RAM. The company will offer a range of packages with the largest providing the most I/O pins. The NX-40, for example, will initially come in a 17mm BGA and 9.5mm fine-pitch BGA with 74 HP I/Os, as well as a 10mm QFN with only 22 HP I/Os and four D-PHY lanes. Lattice expects to qualify the first NX-40 variants for production by 3Q20. The smallest NX-17 will have 20 HP I/Os and four D-PHY lanes in a 3.7mm x 4.1mm wafer-level chip-scale package (WLCSP).

Customers can start CrossLink-NX designs now using Lattice's Radiant 2.0 software tools and soft cores for D-PHY, SGMII, OpenLDI, and a PCIe wrapper. The design software includes editors for device, placement, and timing constraints; timing and netlist analysis; a physical-routing view; a power calculator; and a debugger that works with embedded logic analysis.

CrossLink-NX Gets Smart

To accelerate AI inference, Lattice added DSP blocks to CrossLink-NX. Like those in the ECP5 devices, the 18x18

multipliers handle 16-bit integer (INT16) data for inference. Each DSP block is also configurable as two 9x9 multipliers, doubling throughput for 8-bit integer (INT8) math. To ease customers' AI designs, Lattice offers an RTL and software package called SensAI (see [MPR 5/20/19](#), "Lattice Lessens Power With SensAI"). The package currently supports ECP5 and Ice40 FPGAs but will add CrossLink-NX later this year.

SensAI RTL implements a convolution engine that can handle 3x3 compute kernels using nine multipliers arranged in parallel. The NX-40 can implement six such engines, yielding 108 INT8 multiply-accumulate operations per cycle. (This configuration requires 54 of the 56 available 18-bit DSP units for the multipliers; the adders are implemented in programmable logic.) The DSPs operate at a maximum clock speed of 350MHz, generating a peak rate of 75.6 billion operations per second (GOPS)—enough for small models or low frame rates.

The example design in Figure 1 can combine video bridging and AI inference in a single FPGA. For industrial machine vision, the AI logic can perform object detection or counting, offloading this function from an IoT processor to increase performance or cut power. The FPGA's programmable I/Os can support a proprietary interface on the camera side and standard MIPI or PCIe interfaces on the system side.

Automotive applications can use CrossLink-NX to aggregate up to 14 image sensors into a single D-PHY port, enabling an SoC to handle more cameras than its MIPI ports natively allow. In this case, the FPGA logic can simply multiplex the inputs, or it can stitch frames together horizontally and vertically. Surround-view systems, for example, employ four or more cameras to create a 360-degree bird's-eye view of the car.

Avoiding the Duopoly

Vendors of low-end FPGAs like to compare their products with those from the FPGA market leaders, but Intel and Xilinx invest little in low-end chips. Direct CrossLink-NX competitors include the Intel Cyclone V and Xilinx Artix-7 28nm FPGAs, but they also include products from Microchip and startup Efinix. Compared with CrossLink-NX-40, most competing chips derive from much larger designs, increasing cost and package size.

Efinix, however, started its lineup from low densities, and its first chips are now in production (see [MPR 4/23/18](#), "Efinix Samples Its First FPGAs"). Its T20 has about 20,000 logic cells, similar to the CrossLink-NX-17. Although the chip is built in a SMIC 40nm process, Efinix rates core leakage current at only 6.7mA; NX-17 specifications are unavailable, but static current could be in the same range. Like the Lattice device, the T20 includes two MIPI D-PHY ports with four lanes each. It supports only up to 1.5Gbps per lane, matching the 40nm CrossLink. The startup

	CrossLink-NX-40	CrossLink-NX-17	CrossLink
Logic Cells	39,000	17,000	5,900
Embedded RAM	1,512Kb	432Kb	180Kb
Large RAM	1,024Kb	2,560Kb	None
18x18 Multipliers	56	24	None
INT8 Performance	78.4 GOPS	33.6 GOPS	Not applicable
MIPI D-PHY	2 ports, 8 lanes, 2.5Gbps	2 ports, 8 lanes, 2.5Gbps	2 ports, 8 lanes, 1.5Gbps
PCI Express	Gen2 x1	None	None
HP I/O	Up to 74, 1.5Gbps	Up to 48, 1.5Gbps	22, 1.2Gbps
Process	28nm FD-SOI	28nm FD-SOI	40nm CMOS
Availability	Samples 4Q19	Samples 4Q20 (est)	Production 2017

Table 1. CrossLink device comparison. Relative to the previous generation, the NX family increases logic density while adding large-RAM blocks, multipliers (DSP), and PCIe. (Source: Lattice)

offers a T20-based MIPI development kit aimed squarely at the same designs as CrossLink, but it lacks AI cores to compete with SensAI.

Lattice faces indirect competition from processors that integrate AI accelerators and multiple MIPI ports. Intel's Movidius Myriad X was an early example, and it typically serves as a coprocessor (see [MPR 12/17/18](#), "Intel Gains Myriad Customers"). The edge-AI market is moving quickly, however, with automotive and smart-camera SoCs integrating multiple MIPI ports. Over time, these SoCs will reduce Lattice's available market to designs requiring proprietary camera interfaces.

Refresh Powered by Process

As the first Nexus product, CrossLink-NX delivers a desirable combination of high performance, low power, and small footprint. It's best suited to industrial and IoT video applications that require bridging combined with object detection or counting. Although many of these designs are line powered, some industrial sensors added to existing factories use battery power. Leakage power is important in industrial and other IoT designs with low duty cycles, as it can dominate battery life, so the FD-SOI FPGA is ideal for these systems. We view automotive designs as more opportunistic, since advanced driver-assistance systems are rapidly evolving. CrossLink-NX is ideal if the application requires an FPGA for video bridging, but SoCs with integrated AI accelerators provide greater performance.

Although Lattice withheld a detailed roadmap, it expects to sample the second Nexus product in 1H20 followed by a third in 2H20. CrossLink-NX addresses its largest segment by revenue, comprising industrial and automotive

Price and Availability

CrossLink-NX-40 samples are now available, whereas NX-17 samples are due in 4Q20. Lattice withheld pricing, which we estimate starts around \$15 in 1,000-unit volumes. More information is at www.latticesemi.com/en/Products/FPGAandCPLD/CrossLink-NX.

markets. The company's two other segments generate similar revenue: communications and computing, and consumer and broadcast. For the former, we expect Lattice is developing a larger Nexus device as a follow-on to its ECP5 line, which has up to 84,000 logic cells. For consumer applications, it will likely scale Nexus down to replace its Ice40 line that offers power as low as 10mW (see [MPR 2/23/15](#), "Honey, I Shrunk the FPGA").

There's nothing revolutionary about the first Nexus products—they largely move existing intellectual property to 28nm FD-SOI, reducing power and increasing logic density. We see opportunity, however, for Lattice to expand Nexus in both high-performance and low-power directions by optimizing its use of back bias. Just as important is the company's greater focus on software stacks and soft intellectual property, which is critical to easing customers' FPGA-based designs. For CrossLink-NX, the prime example of this solution strategy is SensAI, bringing support for neural-network frameworks such as Caffe and TensorFlow. If it executes to plan, Lattice should exit 2020 with a refreshed product line reflecting its improved focus. ♦

To subscribe to *Microprocessor Report*, access www.linleygroup.com/mpr or phone us at 408-270-3772.