By: Jeffrey Burt | December 04, 2018

The company unveiled its next-generation programmable networking ASIC, which will start sampling in the first half of next year.

Barefoot Networks officials for the past couple of years have been driving programmability into networks through its Tofino Ethernet ASIC and the P4 programmable language to address the rising demand for more bandwidth and features to address new workloads like artificial intelligence and machine learning.

ASICs for more than two decades have been fixed-function chips, wired to do only one function, according to Ed Doe, vice president of product, business and strategy for Barefoot. However, the trend within hyperscale and enterprise data centers has been toward greater programmability in silicon and domain-specific architectures GPUs and digital signal precessors have grown in importance to help systems better handle particular tasks, and now there are technologies like Tensor processing units to address AI workloads.

Intel and Xilinx also have driven the development of field programmable gate arrays (FPGAs), chips that can be reprogrammed via software. What's been needed is similar programmability in networks, Doe told eWEEK.

“Networks have been relatively stagnant over the last 20 years; we’re still using TCP and IP,” which Doe credited for their longevity and robustness, “but I don’t think many people in 1994 knew what the internet was going to be today. There’s also a lot of new next-generation workloads … with a lot of machine learning, AI, inference [and] training, and the network has to change to really enable those.”

Barefoot has aimed to solve the challenge with Tofino, an Ethernet chip that started appearing in systems last year and has been adopted by such top-tier networking vendors as Cisco Systems (in its Nexus 3400 switch) and Arista Networks (in the 7170 series switches) and hyperscalers like Google, Alibaba and Tencent.

Tofino 2

Now the 5-year-old company is introducing Tofino 2, the next-generation programmable switch chip that is built via a 7-nanometer manufacturing process and will be offered in three different families that will target a broad array of companies and workloads. Like its 16nm predecessor, Tofino 2 leverages Barefoot’s Pisa architecture and uses the P4 programmable language that enables the programming of packet forwarding planes. The company in September rolled out its P4 Studio, a software development environment created to help drive adoption of programmable forwarding planes.

Tofino 2 will start sampling in the first half of next year.

“In networking, there’s really never been this concept of a flexible, programmable domain-specific architecture, where there’s a language, where there’s a compiler and where there’s an architecture that really provided the characteristics specific to networking that provides the flexibility,” Doe said.

Speed and bandwidth are key drivers behind the need for greater programmability in networking, he said. The trend now is toward high-speed networks, up to 100 Gigabit Ethernet and with 400GbE becoming a sizable part of the market by 2020. The Tofino chips will help meet the need for more bandwidth and lower latency, which will increasingly be needed for such workloads as machine learning training and inference. Emerging memory technologies, including storage-class memory (SCM), also are driving the need for larger and more scalable networks.

Having easily programmable networks that can adapt to rapidly changing demands will be crucial going forward, Doe said. He also disputed the idea that greater programmability and flexibility has to come with tradeoffs around performance, cost and features. In fact, a switch powered by a Tofino chip delivered 21 percent better performance, 53 percent better performance-per-watt and two to 20

times the performance in such areas as routing than the same switch with a fixed-function ASIC, he said.

Tofino 2 will offer up to 256 10, 25 and 50GbE ports on a single chip, or 32 400GbE ports. It can be deployed on everything from a top-of-rack switch to a server provider router to a switch appliance, delivers twice the bandwidth (12.4 Tb/s) as its predecessor and is built using a modular chiplet architecture with silicon that can be upgraded up to 100G SerDes and silicon photonics.

Three Tofino Chip Families

The new chip also will come in three families (the first Tofino chip had a single family). The M-series is aimed the same mainstream enterprise and data center spine-and-leaf architecture that the initial Tofino ASIC targeted. The U-series is aimed at high-capacity and high-feature environments, such as 5G and edge computing for service providers, next-generation storage and hyperscalers with machine learning and deep learning fabrics, load balancers and firewalls.

The H-series is optimized with power and efficiency for scaled-down hyperscale environments that have less complexity and need low latency.

The new chip families will help Barefoot expand its reach in hyperscale environments as well as among telecommunications companies and service providers, Doe said. The introduction of new features and the support of the P4 programming language are key differentiators for Barefoot, he said, comparing the situation with that of GPU makers Nvidia and Advanced Micro Devices. Before the development of such technologies as CUDA, GPUs were used primarily for graphics jobs in PCs.

“But when they opened it up with the programming language, it allowed them to expand to a lot of these new nascent markets that might have been really nichy, like virtual reality and AR [augmented reality], or HPC [high-performance computing] and big data,” Doe said. “But these markets were exploratory markets because people were able to start using that same engine for these new applications.”