

Enhanced Visual Intelligence at the Network Edge

Intel® Movidius™ Myriad™ X VPU with Neural Compute Engine



Take your imaging, computer vision and machine intelligence applications into network edge devices with the newest Movidius family of vision processing units (VPUs) by Intel.

Industry Leading Performance at Ultra-Low Power

Intel's Myriad X third generation VPU delivers class-leading performance in computer vision and deep neural network inferencing applications. As the newest member of the Movidius VPU family known for ultra-low power consumption, the Myriad X VPU is capable of delivering a total performance of over 4 trillion operations per second (TOPS).¹ With new performance enhancements, the Myriad X VPU is a power efficient solution that brings advanced vision and artificial intelligence applications to devices such as drones, smart cameras, smart home, security, VR/AR headsets, and 360 cameras.

New Generation of Deep Neural Network Performance

Intel has introduced an entirely new deep neural network processing unit into the Myriad X VPU architecture: the Neural Compute Engine. Specifically designed to run deep neural networks at high speed and low power, the Neural Compute Engine enables the Myriad X VPU to reach over 1 TOPS of compute performance on deep neural network inferences.² The Neural Compute Engine is integrated as part of the power efficient Movidius VPU architecture which minimizes power by reducing data movement on-chip. While the Myriad 2 VPU has provided superior deep neural network support at low power, the Myriad X VPU can now reach 10X higher performance for applications requiring multiple neural networks running simultaneously.²

Customizable Imaging & Vision Pipelines

The Movidius family of VPUs have always provided a unique, flexible architecture for image processing, computer vision, and deep neural networks. The architecture provides a modular approach to configuring imaging and vision workloads because it combines a set of imaging and vision hardware accelerators, such as stereo depth or the Neural Compute Engine, with an array of C-programmable VLIW vector processors, all accessing a common on-chip memory. This approach enables world-class image signal processing (ISP) without the need to make trips to memory for best power efficiency, in addition to interleaved computer vision and deep neural network inference application pipelines, all with a dataflow methodology that reduces power by minimizing data movement. Movidius VPUs deliver an optimal balance between programmability and performance at low power.

Support for 8 HD Sensors and 4K Encoding

The Myriad X VPU features 16 MIPI lanes, which supports up to 8 HD resolution RGB sensors to be connected directly. The high-throughput inline ISP ensures streams are processed at high speeds, while new hardware encoders provide support for 4K resolutions at both 30 Hz (H.264/H.265) and 60 Hz (M/JPEG) frame rates. Other featured interfaces include USB 3.1 and PCI-E Gen 3.

Software Development Kit (SDK) and Tools

The Myriad X VPU ships with a rich SDK that contains all of the software development frameworks, tools, drivers and libraries to implement custom imaging, vision and


deep learning applications on Myriad X VPU. The SDK also includes a specialized FLIC framework with a plug-in approach to developing application pipelines including image processing, computer vision, and deep learning. This framework helps developers focus on the processing, leaving dataflow optimization to the tools. For deep neural network development, the SDK includes a neural network compiler that enables developers to rapidly port neural networks from common frameworks such as Caffe* and Tensorflow* with an automated conversion and optimization tool that maximizes performance while retaining network model accuracy.

Where to Get More Information

For more information, visit www.movidius.com/MyriadX

Movidius® Myriad™ X VPU at a Glance

FEATURES	BENEFITS
Neural Compute Engine	With this dedicated on-chip accelerator for deep neural networks, the Myriad X VPU delivers over 1 trillion operations per second of DNN inferencing performance. ² Run deep neural networks in real time at the edge without compromising on power consumption or accuracy.
16 Programmable 128-bit VLIW Vector Processors	Run multiple concurrent imaging and vision application pipelines with the flexibility of 16 vector processors optimized for computer vision workloads.
16 Configurable MIPI Lanes	Connect up to 8 HD resolution RGB cameras directly to the Myriad X VPU with support for up to 700 million pixels per second of image signal processing throughput.
Enhanced Vision Accelerators	Utilize over 20 hardware accelerators to perform tasks such as optical flow and stereo depth without introducing additional compute overhead. For example, the new stereo depth accelerator can simultaneously process 6 camera inputs (3 stereo pairs) each running 720p resolution at 60 Hz frame rate.
2.5 MB of Homogenous On-Chip Memory	The centralized on-chip memory architecture allows for up to 400 GB/sec of internal bandwidth, minimizing latency and reducing power consumption by minimizing off-chip data transfer.
2 chip packages offered	MA2085: No memory in-package; interfaces to external memory MA2485: 4 Gbit LPDDR4 memory in-package



¹ Overall performance is the architectural calculation based on maximum performance of operations-per-second over all available compute units. Application performance varies based on the application.

² Maximum performance based on peak floating-point computational throughput of Neural Compute Engine. Actual results on deep neural networks may achieve less than peak throughput.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Intel, Movidius, and Myriad are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. * Other names and brands may be claimed as the property of others.