SYNOPSYS®
Silicon to Software™

Taking Configurability and Programmability to the Next Level for Embedded Systems Processing

Dr. Yankin Tanurhan
Vice President of Engineering, SG
Synopsys Inc.

July 2016

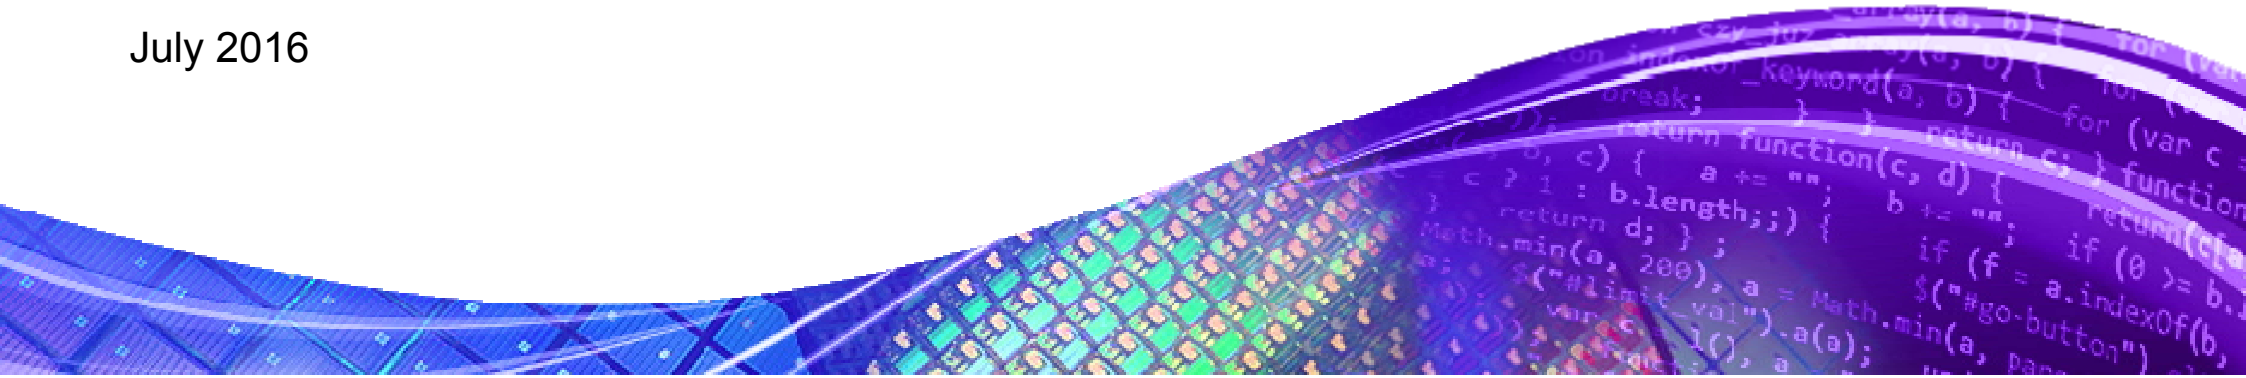

Agenda

Machine vision

Scalable EV Processors

Convolution Neural Network Engine (CNN)

Vision programming tools

Host Integration

Conclusions

Embedded Vision is Coming Fast

- Embedded Vision is the use of computer vision in embedded systems to interpret meaning from images or video
- In cars to improve safety
- Surveillance for detection and tracking
- In industrial automation to improve quality and control
- Estimated $300B+ market in 2020, 35% CAGR

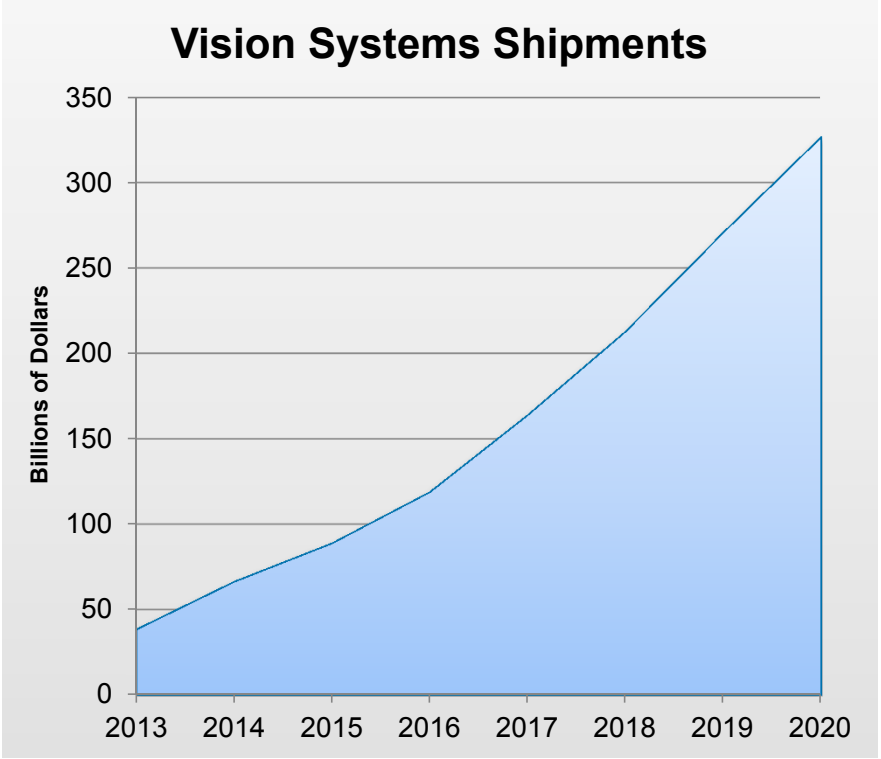

Sources: ABI Research, Insight Media, Transparency Market Research, Markets And Markets, Synopsys

Video Surveillance Market

Home Surveillance, Retail, Healthcare, Security (Airports, Govt, Banks, Casinos)

- Global IP Video Surveillance Market expected to grow at CAGR of 37.3% from 2012-20
- Demand driven by
  - Growing installations of IP cameras
  - Need for surveillance cameras with better video quality

FIG. 1 Global video surveillance and VSaaS market size and forecast, 2011 – 2019 (USD billion)

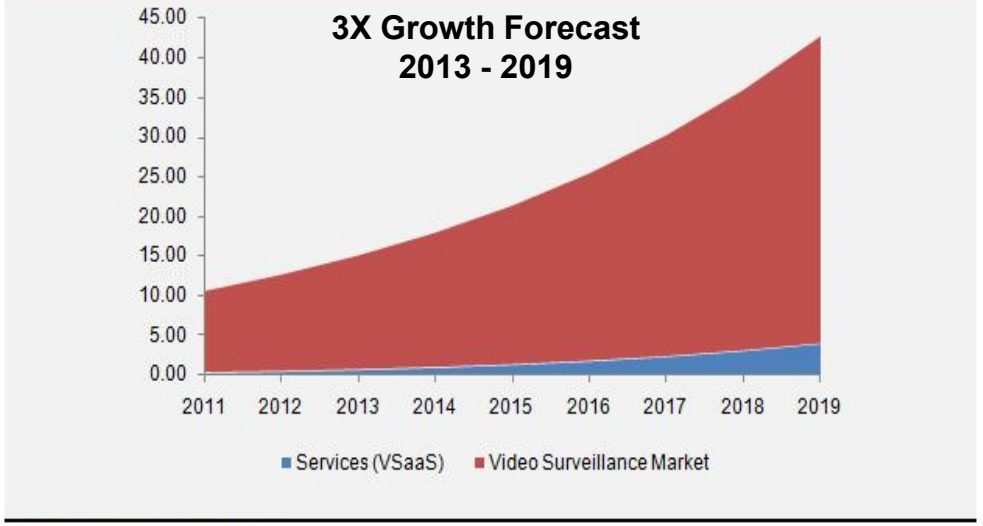

Source: Transparency Market Research

http://www.alliedmarketresearch.com/IP-video-surveillance-VSaaS-market

ADAS Vision Market

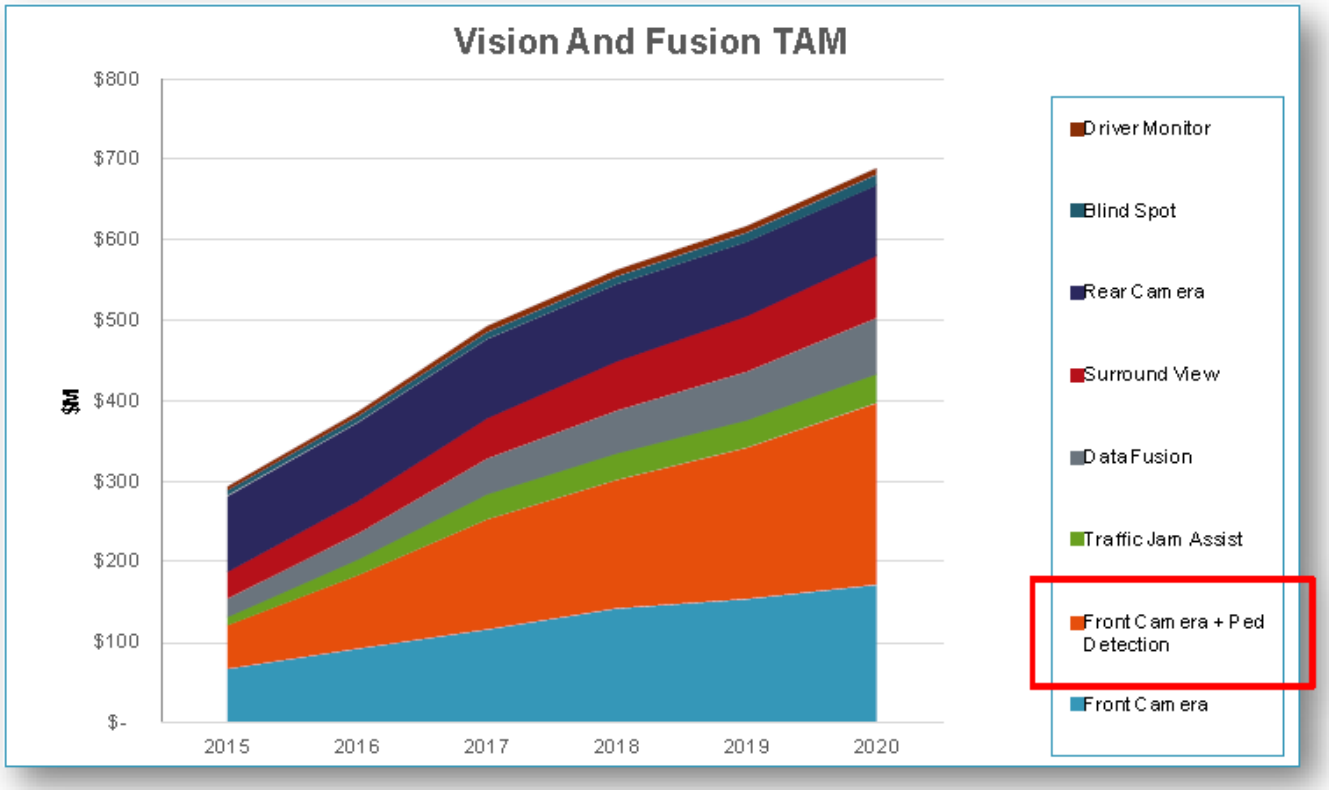

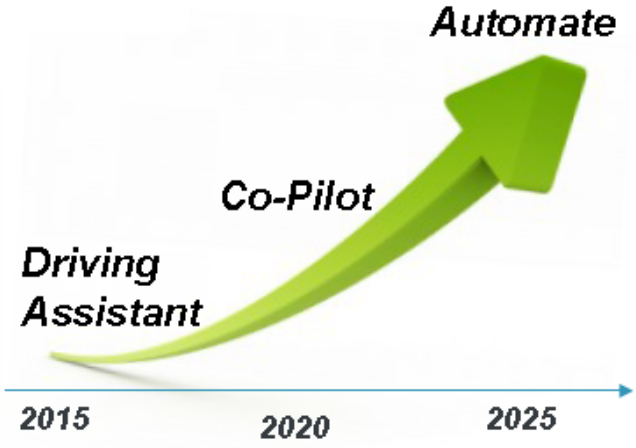

8.2 million *autonomous* cars by 2030

Market: ***20% CAGR,*** driven by safety deployment

Source: IHS 2015

SYNOPSYS®

Scalable Embedded Vision Processors

Scalable Embedded Vision Processors

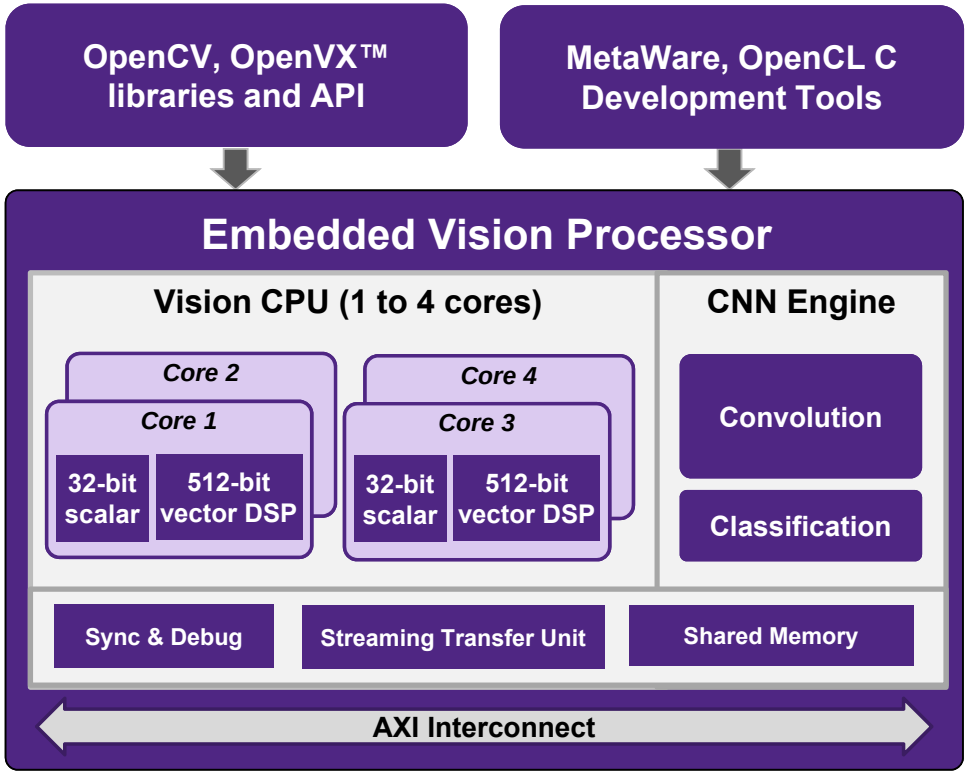

- Highly integrated and configurable
  - Configurable scalar, vector DSP and convolutional neural network (CNN) architecture
  - Supports 1080p - 4K vision streams
- User scalable for optimum performance
  - 1 to 4 Vision CPU cores
  - Programmable CNN engine
- State-of-the-art performance-efficiency
- High productivity toolset
  - OpenCV, OpenVX, OpenCL C

SYNOPSYS

High-Performance Vision CPU

Fast, Real-time Vision Processing

- Implemented with 1 to 4 Vision CPU cores to run full range of vision algorithms up to 4K resolutions
- 32-bit scalar for easy system integration
- Configurable vector DSP for high-end vision processing
- Cross-bar implements scatter/gather delivering higher performance with increased efficiency
- OpenCL C optimized instruction set for high programming productivity with reduced cycle count and lower power
- Parallel operation with the CNN Engine increasing efficiency and throughput

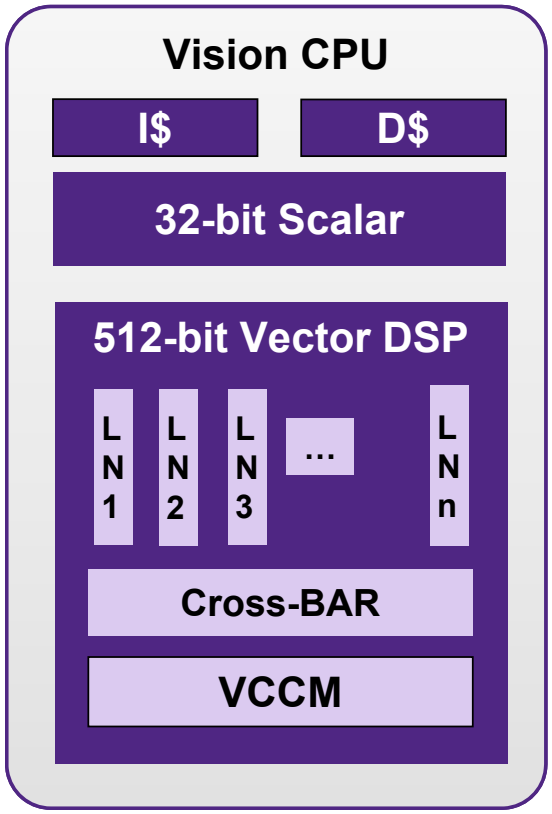

Multi-Core Scalability to Increase Performance

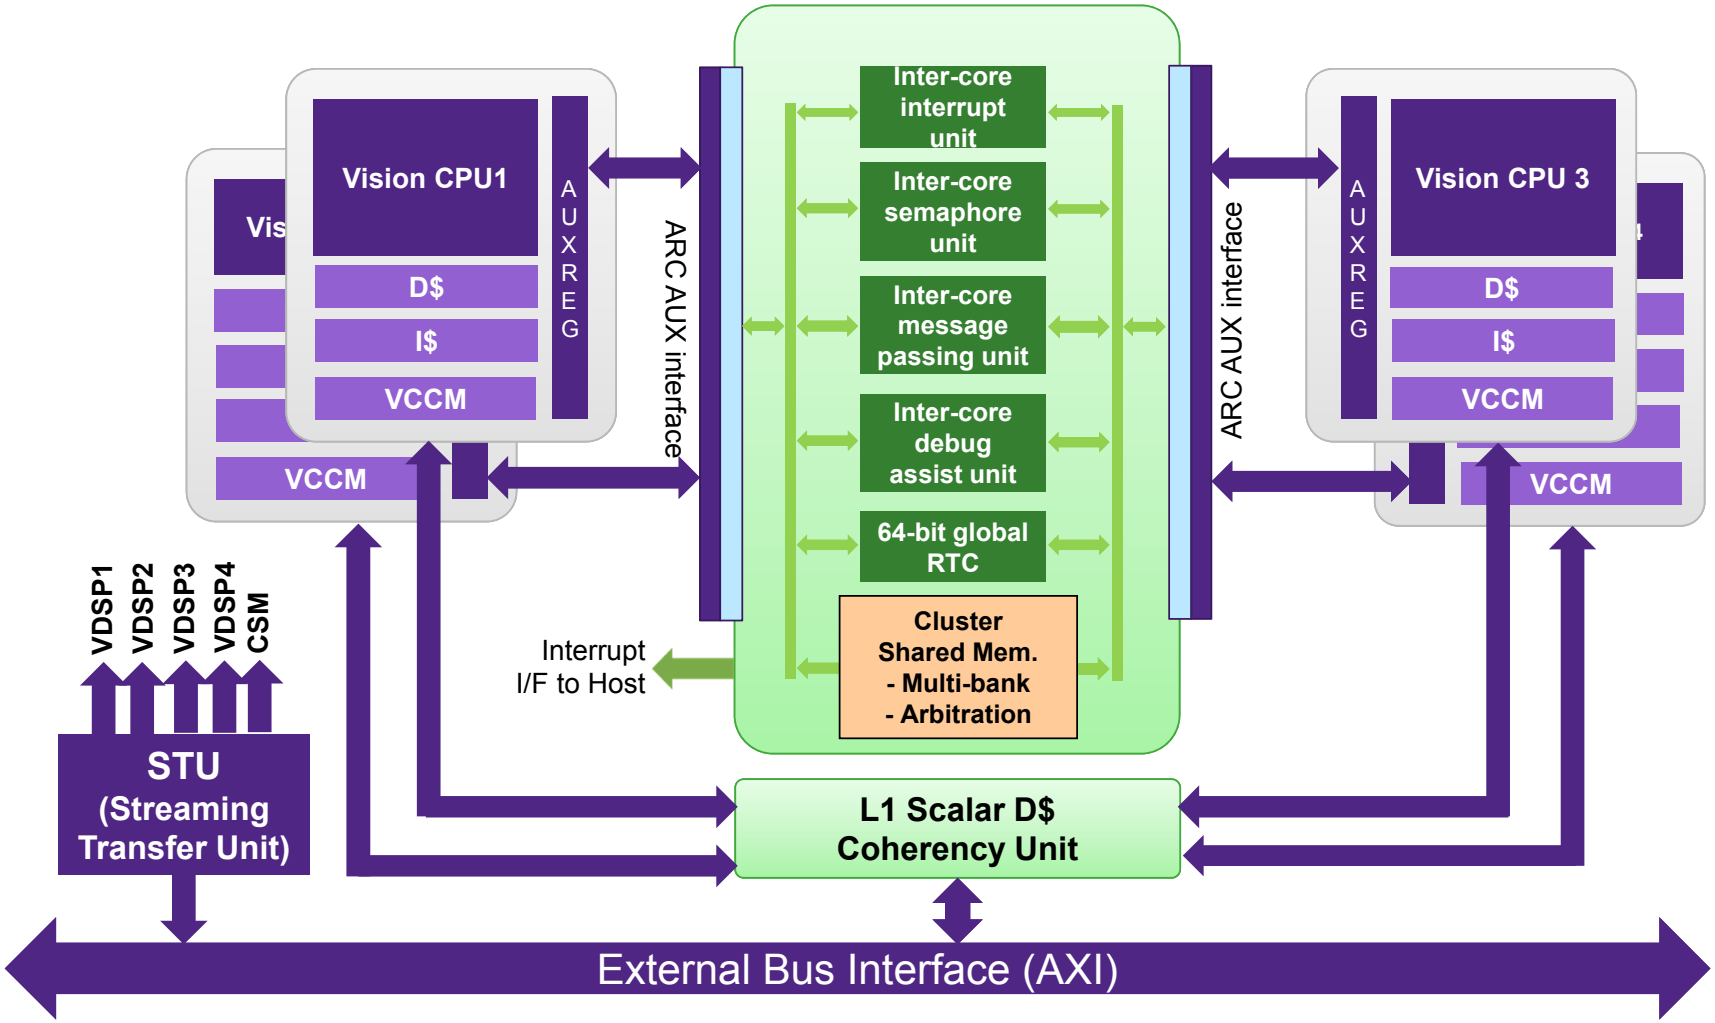

Efficient Frame Data Access

- Streaming Transfer Unit (DMA)
- Tightly integrated with VDSP vector memory (VCCM)
  - Capable of addressing the memory as Private Banks or Local Memory
  - Low overhead:
    - Transfers can be setup in VCCM memory with vector instructions
    - Event-based low-overhead synchronization (no ITs required)
  - Physical channels (up to 4) can be shared between VDSP cores to serve multiple private banks in parallel
- Support for 1D and 2D (image) transfers for vision applications
- I/O coherency: Configurable I/O-coherent regions

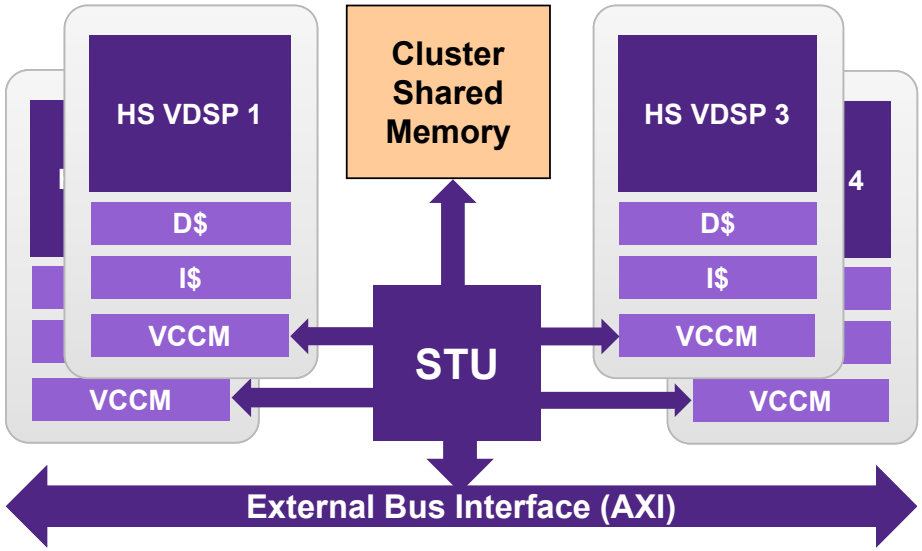

Convolution Neural Network Engine

CNN is Accurate and Efficient

- Convolutional Neural Networks (CNN)
  - Deep learning approach outperforms other vision algorithms
  - CNNs attempt to replicate how the brain sees
  - Efficiently recognizes objects directly from pixel images with minimal pre-processing
- CNN usage for Vision
  - Image classification, search similar images
  - Object detection, classification & localization
    - Any type of object(s), depending on training phase
  - Face recognition, visual attention, facial expression recognition
  - Gesture recognition / hand tracking
  - Scene recognition and labelling, semantic segmentation

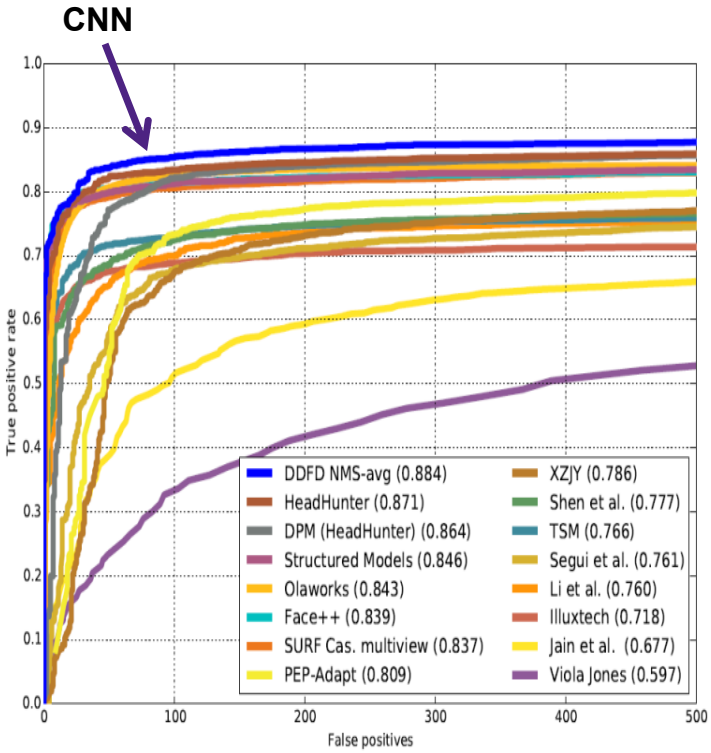

User Benefits of a Dedicated CNN Engine

- Scalable high performance
  - Over 800 MAC/cycle
- PPA competitive with H/W-based CNNs
  - Similar area density (GMAC/s/mm^2) as H/W
  - Over 1000 GOP/s/W
- Low external memory bandwidth
  - Close to image pixel bandwidth <800 MB/s
- Highest productivity and flexibility
  - Automatic CNN graph mapping
- Closely coupled with leading edge EV6x processor
  - Unique combination of scalar + vector processing + CNN
  - Integrated into a single, standard programming model

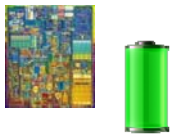

OpenVX

SYNOPSYS

High-Performance CNN Engine

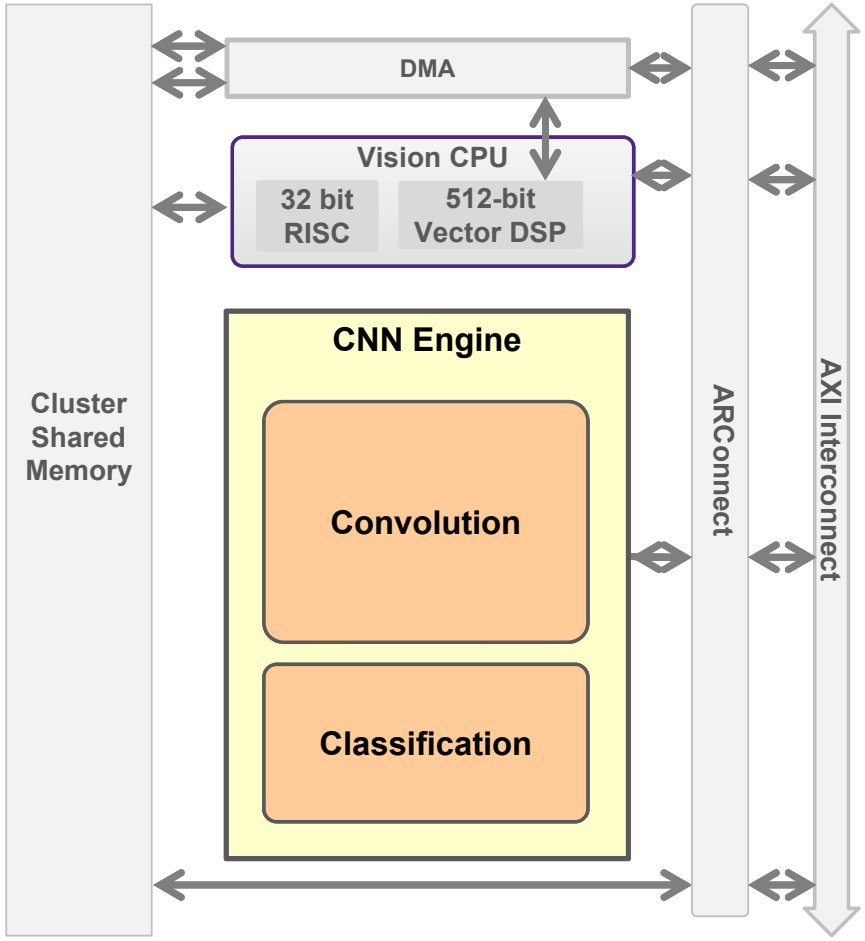

- Dedicated CNN Engine delivers higher performance than competitive solutions
- Programmable to support full range of fixed point CNN graphs
- State-of-the-art power-efficiency >1000 GOPS/W
- Supports resolutions to 4K
- Real-time, high quality image classification, object recognition, semantic segmentation
- Parallel operation with Vision CPU increasing efficiency and throughput

SYNOPSYS

CNN Scene segmentation benchmark

- Application is to perform semantic labeling of full image
  - Define sky, forest, buildings
  - Define border of road (w/o lane markers)
  - Used typically in path finding applications for autonomous driving
  - Unlike object detection, there is no Region-of-Interest pre-calculation
    - Need to analyze entire image
- CNN graph characteristics
  - Video at 1080p, 30 fps
  - 5 channels (3 channels for colors, 2 application-specific channels)
  - Image bandwidth requirements:
    311 MB/s (for 8 bit pixels), or 622 MB/s (for 16 bit pixels)
  - Compute requirements: Over 750 GMAC/s
  - Weights: Over 100K values

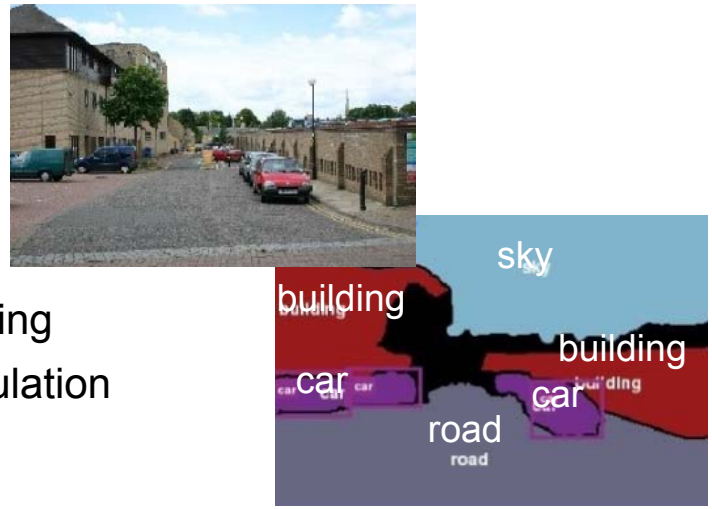

Vision Programming Tools

High Productivity Tools Increase Flexibility

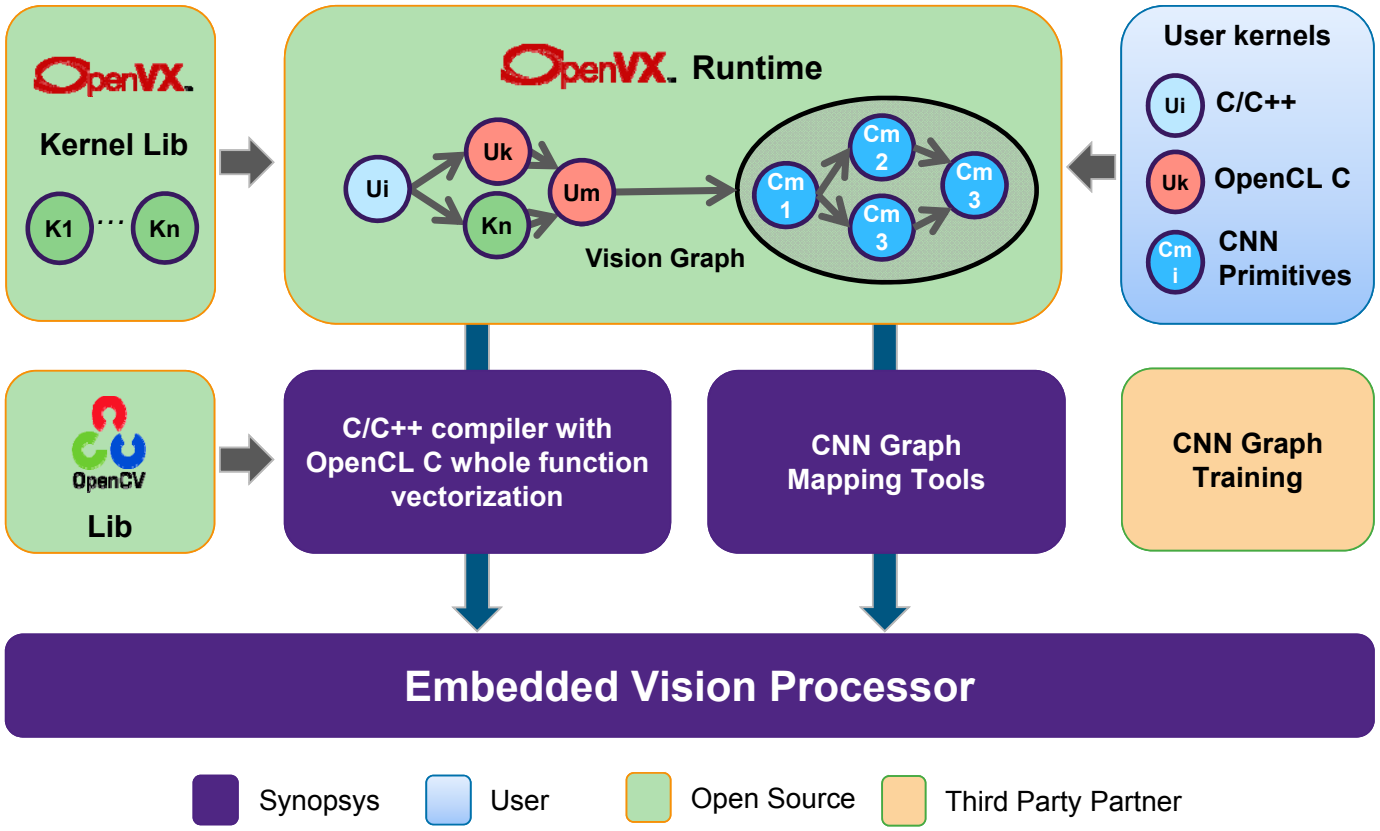

- OpenVX eases vision graph development
- OpenCV open source library of 2500 vision algorithms helps build vision applications
- MetaWare C/C++ compiler delivers optimize program coding
- OpenCL C instructions with whole function vectorization simplifies DSP programming
- CNN graph mapping tools automate CNN programming

SYNOPSYS

FPGA-Based Development System

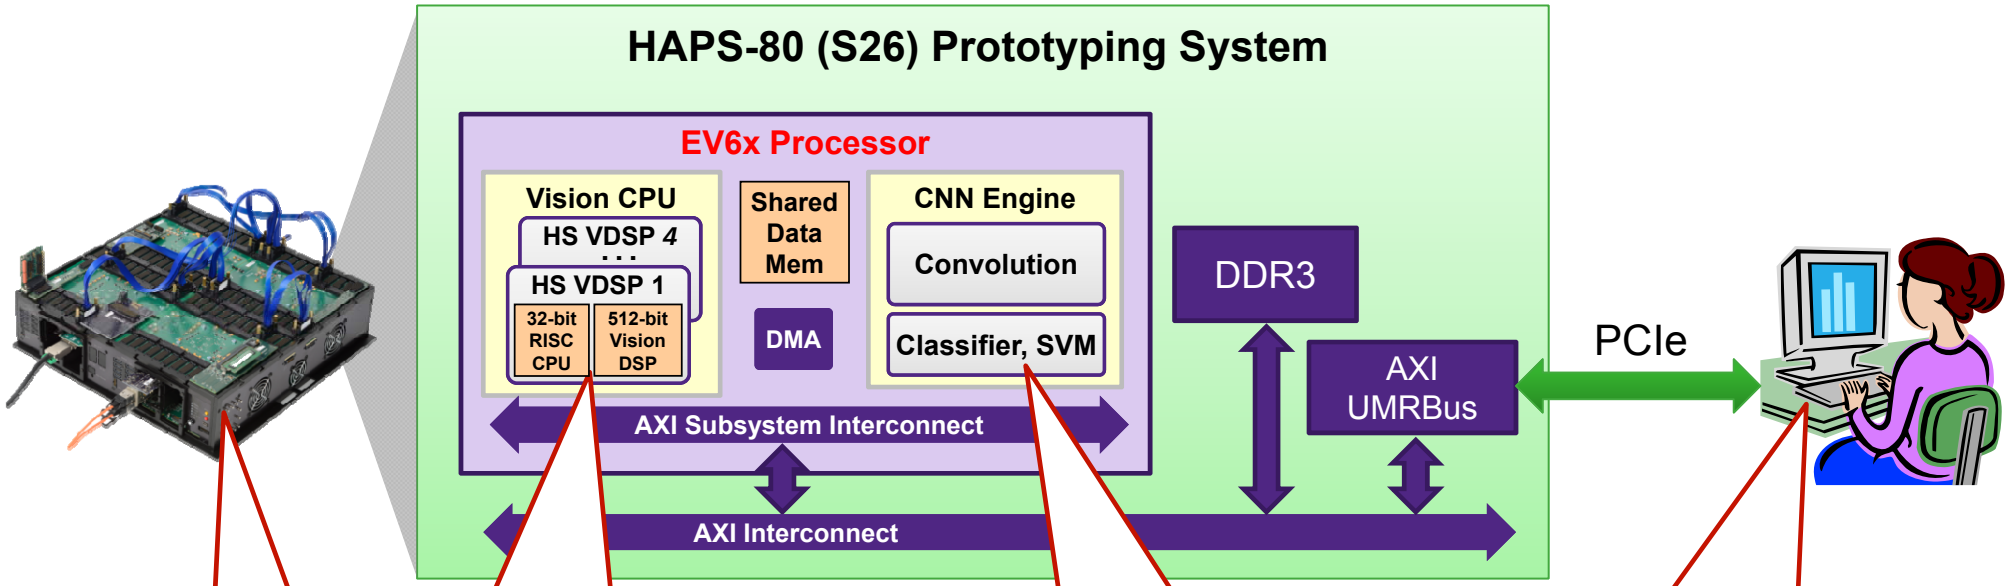

HAPS-80

EV6x runs at up to 100 MHz

Vision CPU

- Scalar & vector processing
- 1 – 4 HS VDSP Cores
- Supports HD to 4K streams
- Host offload or standalone operation

CNN Engine

- Object detection, classification, localization
- Fully programmable
- >1000 GMAC/s/W

Software Development

- High productivity tool set
- OpenVX runtime & kernels
- OpenCL C vectorizing compiler
- MetaWare C/C++ compiler

SYNOPSYS

EV Processor SoC Virtual Development Kit

- Host integration demonstrator
  - VDK with ARC HS38 host running Linux O/S
  - Linked to EV processor model
  - Low-level mechanisms provided to support customer-specific host O/S driver development
  - Simple host-side demonstrator application: e.g., send frame for object detection processing, receive detection points

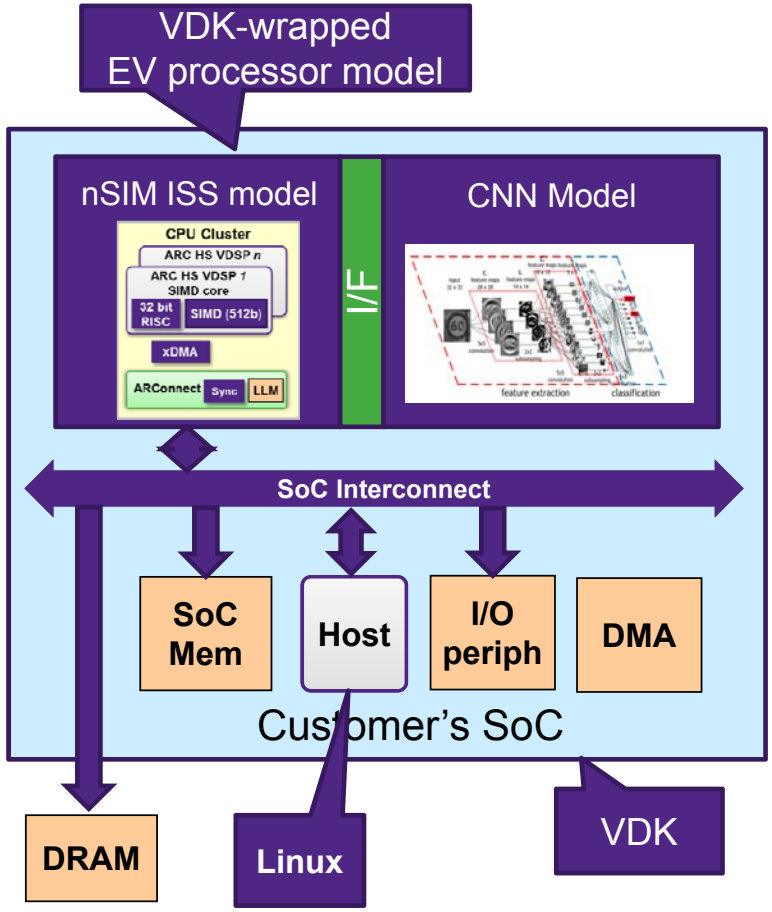

Host Integration

Designed for SoC Integration

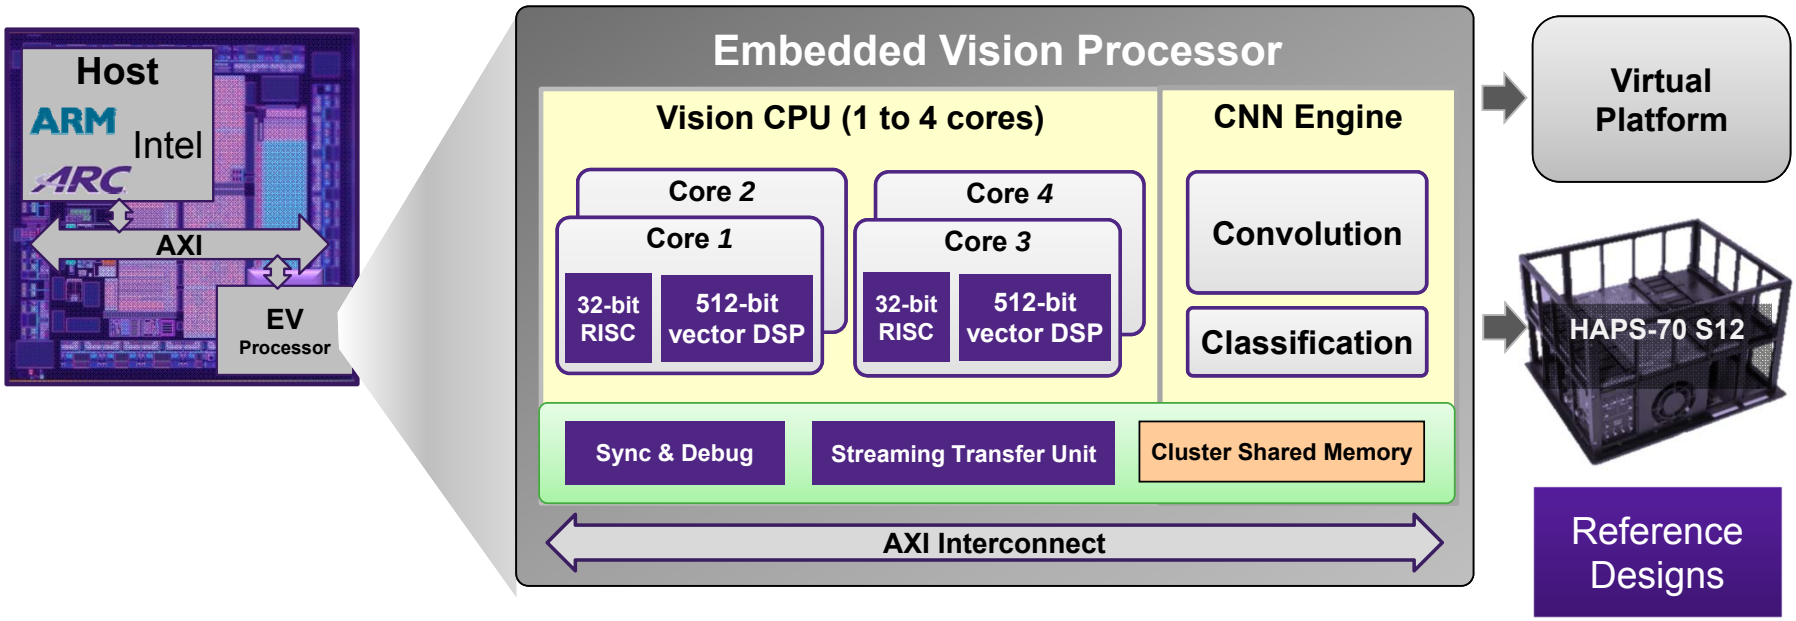

- Must be designed for easy integration with a host processor
  - Most designs will include host

SYNOPSYS

Host to EV6 Communication

As implemented in the EV VDK

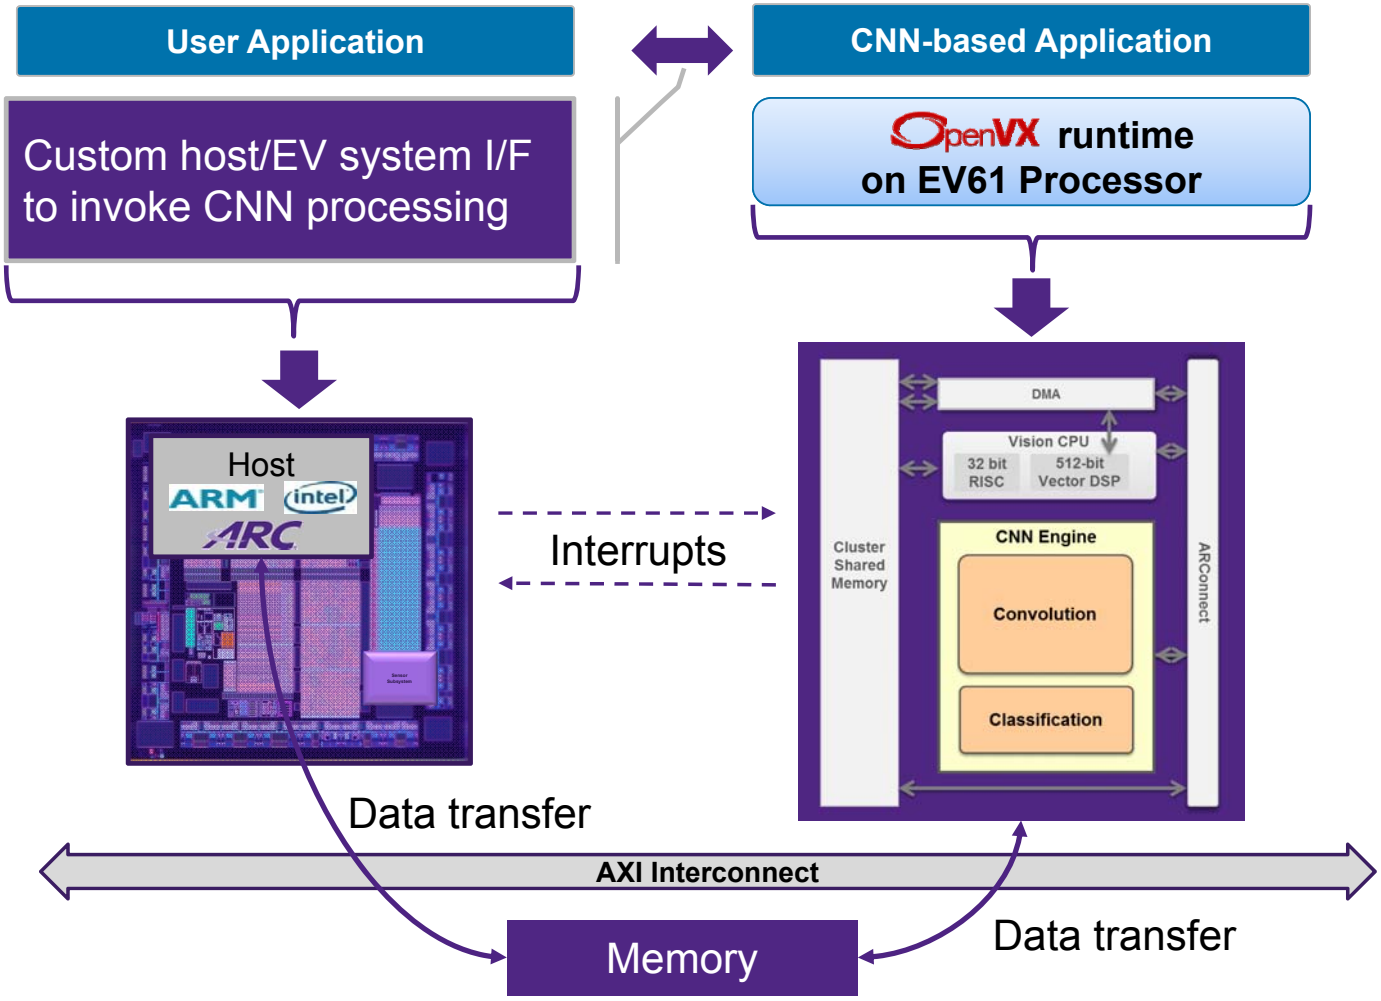

SYNOPSYS

Conclusions

Conclusions

- Strong growth in vision applications
  - Surveillance, ADAS, IoT, Industrial, gesture recognition, object detection
  - Embedded vision algorithm innovation, highly dynamic
  - Requirements for configurability and progammability
- Vision Processors integrate HW & SW for optimized vision apps
  - Flexible, integrated solution (scalar + vector DSP + CNN)
  - Dedicated CNN engine offers higher performance and best efficiency
  - Low-power, high-performance, flexible programming
  - Fast time-to-market with high-level OpenVX and C programming tools
- High productivity tool set
  - OpenVX, OpenCV, OpenCL C, Optimized C/C++ MetaWare compiler

SYNOPSYS®
Silicon to Software™

Thank You