MICROCHIP

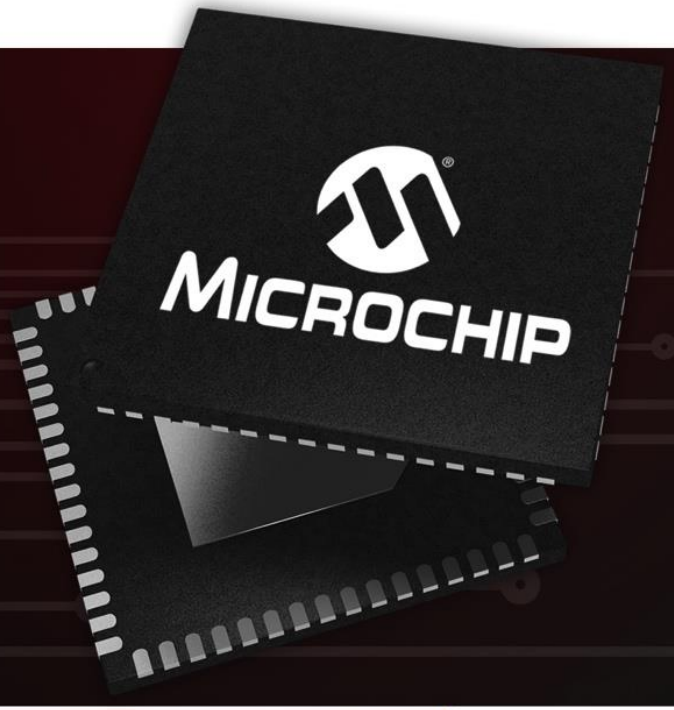

A Leading Provider of Microcontroller, Security, Mixed-Signal, Analog & Flash-IP Solutions



SEV and Machine Learning Inference

May 10, 2019

Machine Learning Inference in Embedded Vision

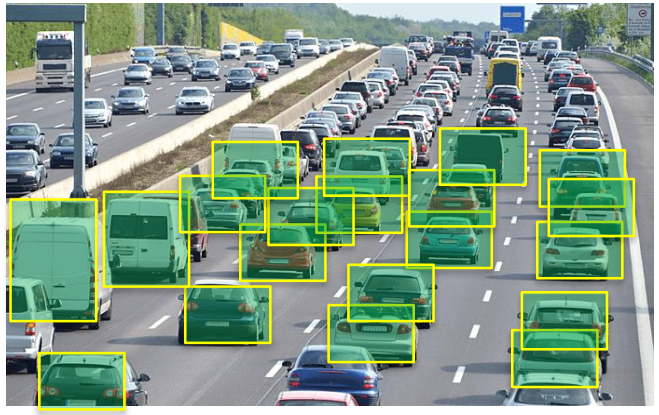

Traffic Monitoring

Machine Vision

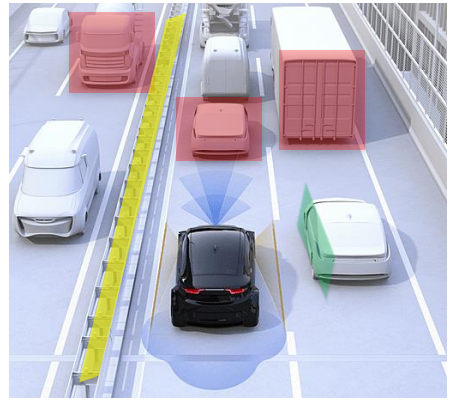

Autonomous Driving

Surveillance

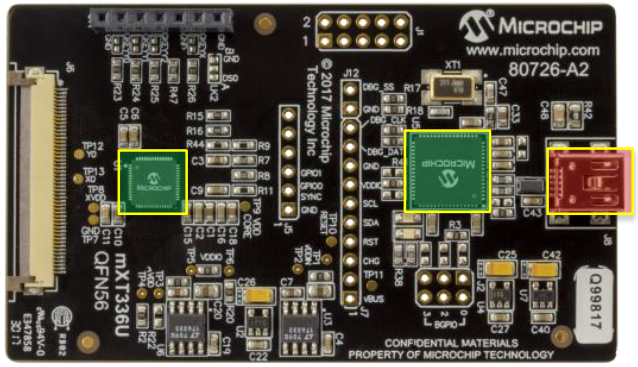

PCB Inspection

Drone Vision

Smart, Embedded Vision: Microsemi Solution

- **Most integrated solution**
  - Imaging / video signal processing
  - Deep learning inference
  - Integrate custom block (e.g. motor control)

5K LE to 500K LE FPGA devices

Video dev kit

- **2x to 4x GOPS / Watt vs other FPGAs**
  - Optimized 8b DOTP block
  - Low power FPGA fabric
  - Optimized implementation

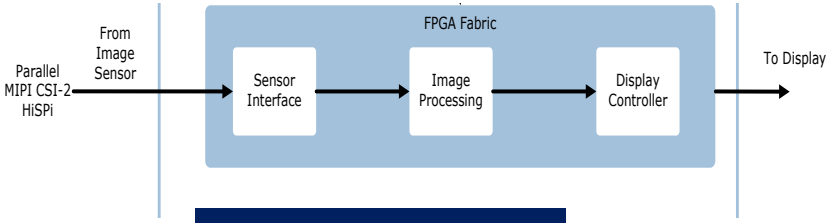

IP & Solutions

- **Most flexible platform**
  - FPGA enables flexibility (e.g. interface)
  - Product family with multiple devices enables resources vs cost trade-off

- **Most secure and reliable**
  - Built in tamper and clone resistance
  - SEU immune FPGA
  - Extended temp range

The Deep Learning Construct

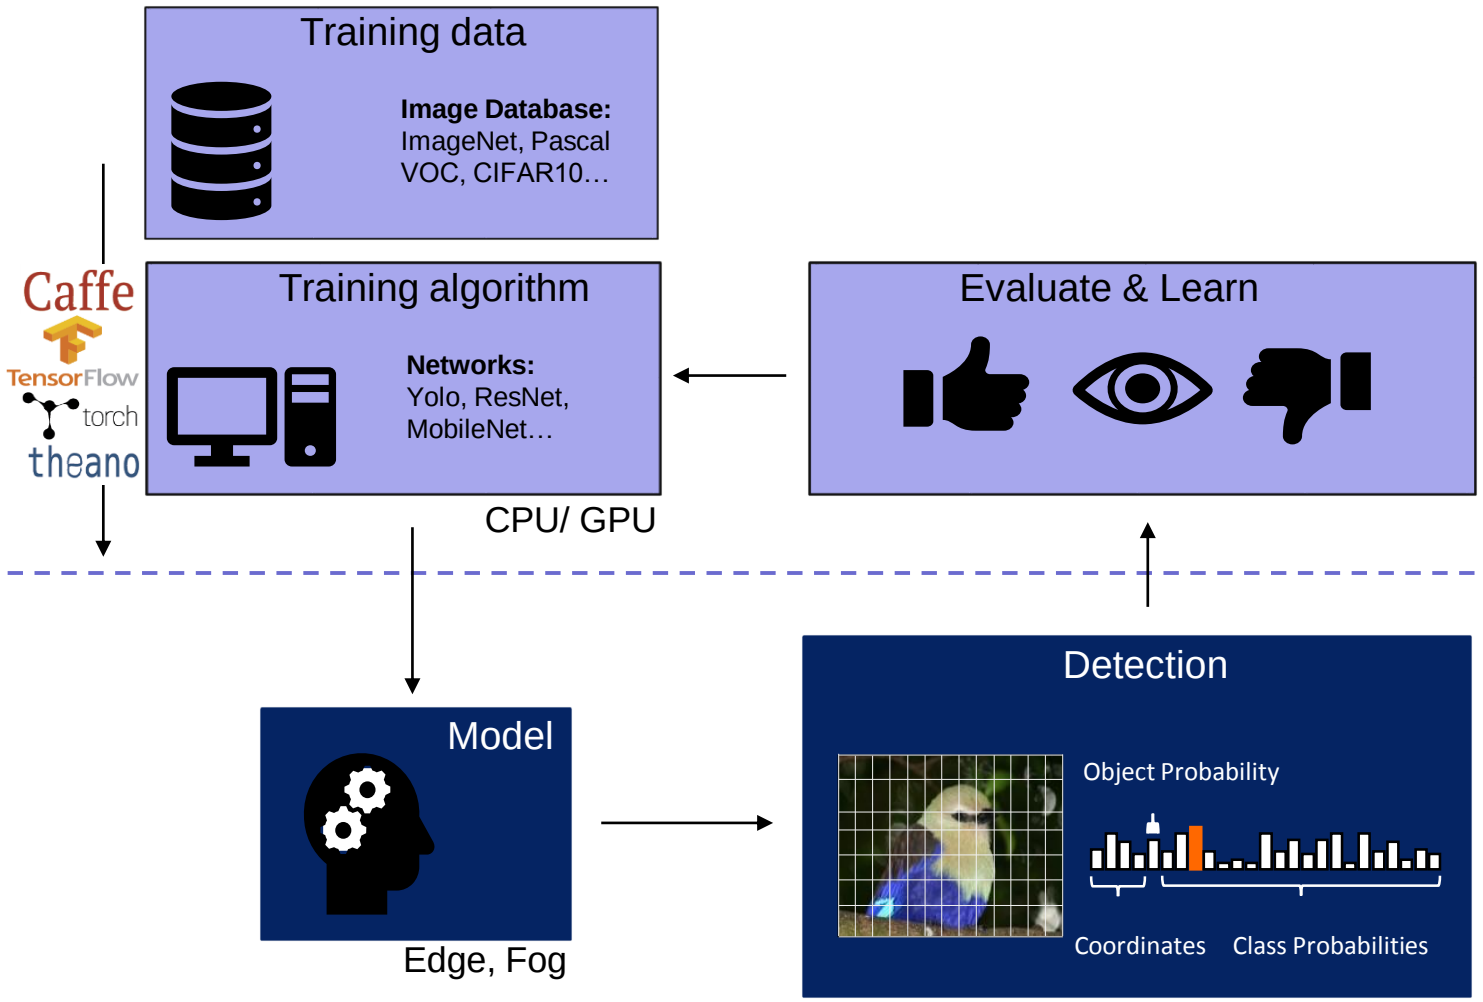

Network Training

- In the data center
- Large compute capacity
- No power or space constraints

Training

Inference

Inference

- Need Low Latency
- Power & space constrained
- Requires security and reliability

FPGAs are Ideal for Inference

1. Large DSP compute capacity (GOPS)

vs. MCU / MPU / CPU

2. Low power dissipation

vs. CPU / GPU

3. Field upgrade

vs. ASIC

4. Integrate more functions

Video, Connectivity, Security…

Example Neural Network

$$o_j = a_1*w_{11} + a_1*w_{12} + a_1*w_{13} + \ldots a_2*w_{21} + a_2*w_{22} + a_2*w_{23} + \ldots$$

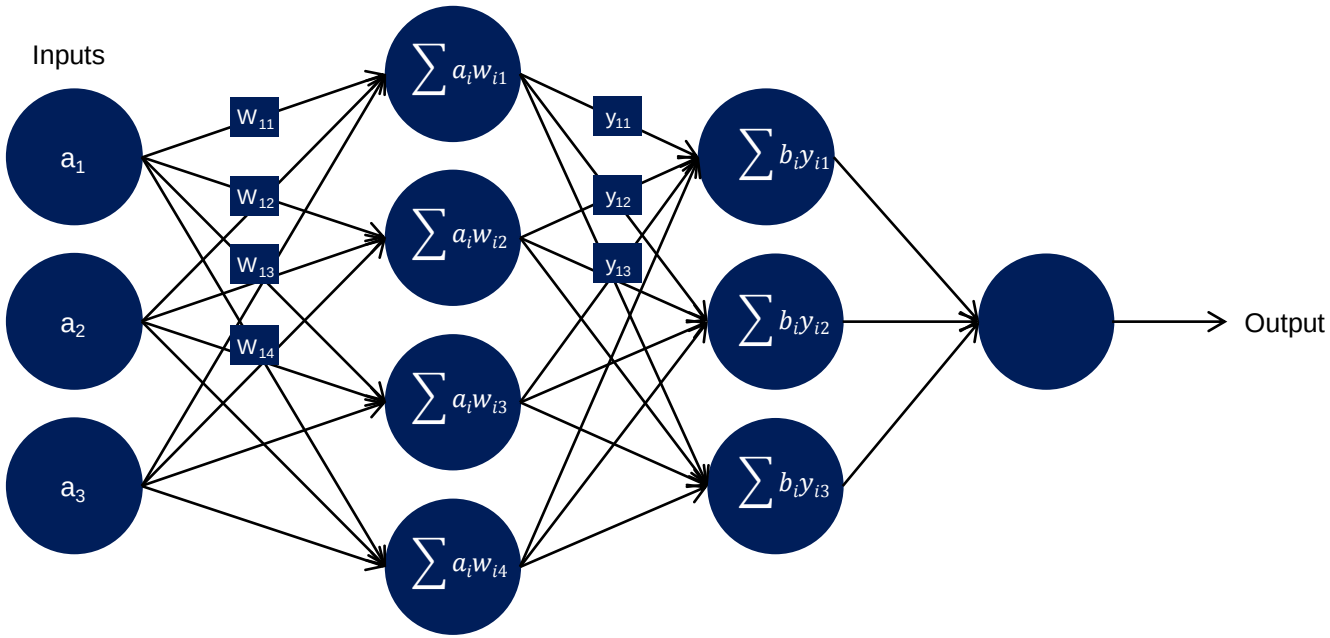

Inputs

$$\begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}$$

Layer 1 Weights

$$\begin{bmatrix} w_{11} & w_{12} & w_{13} & w_{14} \\ w_{21} & w_{22} & w_{23} & w_{24} \\ w_{31} & w_{32} & w_{33} & w_{34} \\ w_{41} & w_{42} & w_{44} & w_{44} \end{bmatrix}$$

Sum-product

$$\sum a_i w_{i1} = b_1$$
$$\sum a_i w_{i2} = b_2$$
$$\sum a_i w_{i3} = b_3$$
$$\sum a_i w_{i4} = b_4$$

Layer 2 Weights

$$\begin{bmatrix} y_{11} & y_{12} & y_{13} \\ y_{21} & y_{22} & y_{23} \\ y_{31} & y_{32} & y_{33} \end{bmatrix}$$

Summary: 2x to 4x GOPS / Watt vs. Competitive SRAM FPGA

INT8 Optimized DSP Block

- Up to 25% higher capacity per math block
- Up to 1480 math blocks
- Up to 1.5 TOPS capacity

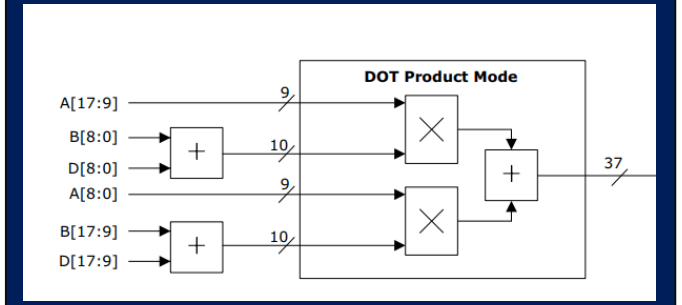

~ 1.25x

CDL Network Optimizing Compile

- Optimized implementation for each network – no shim layer
- Ability to optimize power vs resources vs performance
- Push button implementation

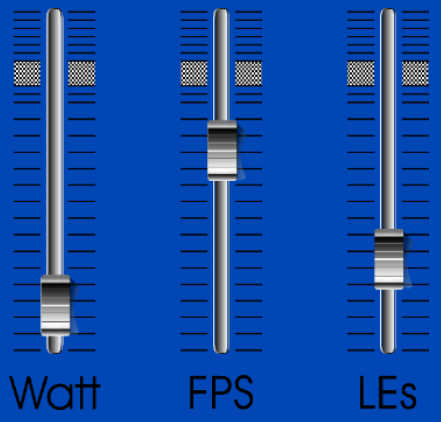

~1.5x

Power Optimized FPGA

- Flash based FPGA reducing static power by 90%+
- Power optimized SERDES
- Total: up to 50% lower total power consumption

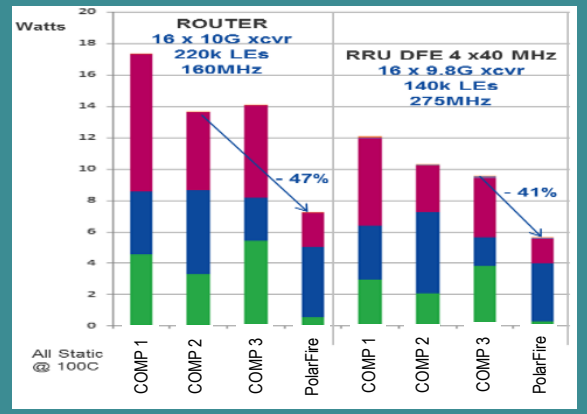

~ 2x

GOPS / Watt compares to equivalent resource FPGA from other vendors; by measuring total power consumption at 100C junction temperature

Resnet v2: Performance Expectations

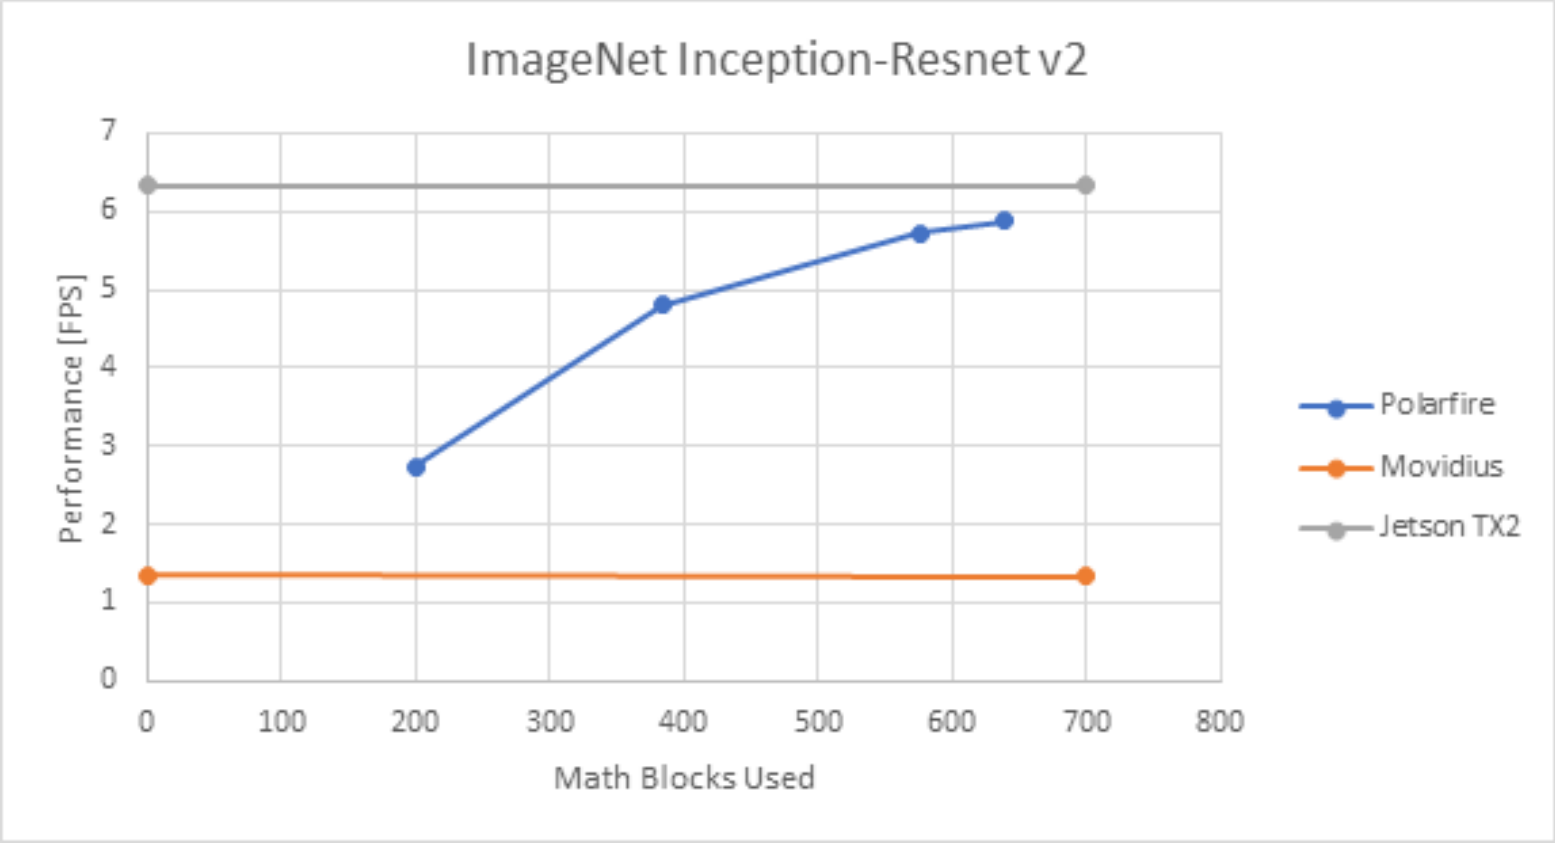

As applied to the ImageNet dataset with 1000 output

PolarFire results based on CDL Design Space Exploration

The results for the Jetson TX2 was obtained from https://github.com/NVIDIA-AI-IOT/tf_trt_models/blob/master/README.md
The Movidius results from https://github.com/movidius/ncappzoo/blob/master/tensorflow/inception_resnet_v2/README.md.

Design Space Exploration enables Scalability

Input image shape	416 x 416
Number of convolutional layers	9
GOPs (MULACC)	7

Up to 150KLE
Up to 240 Math Blocks

Scalable Platform

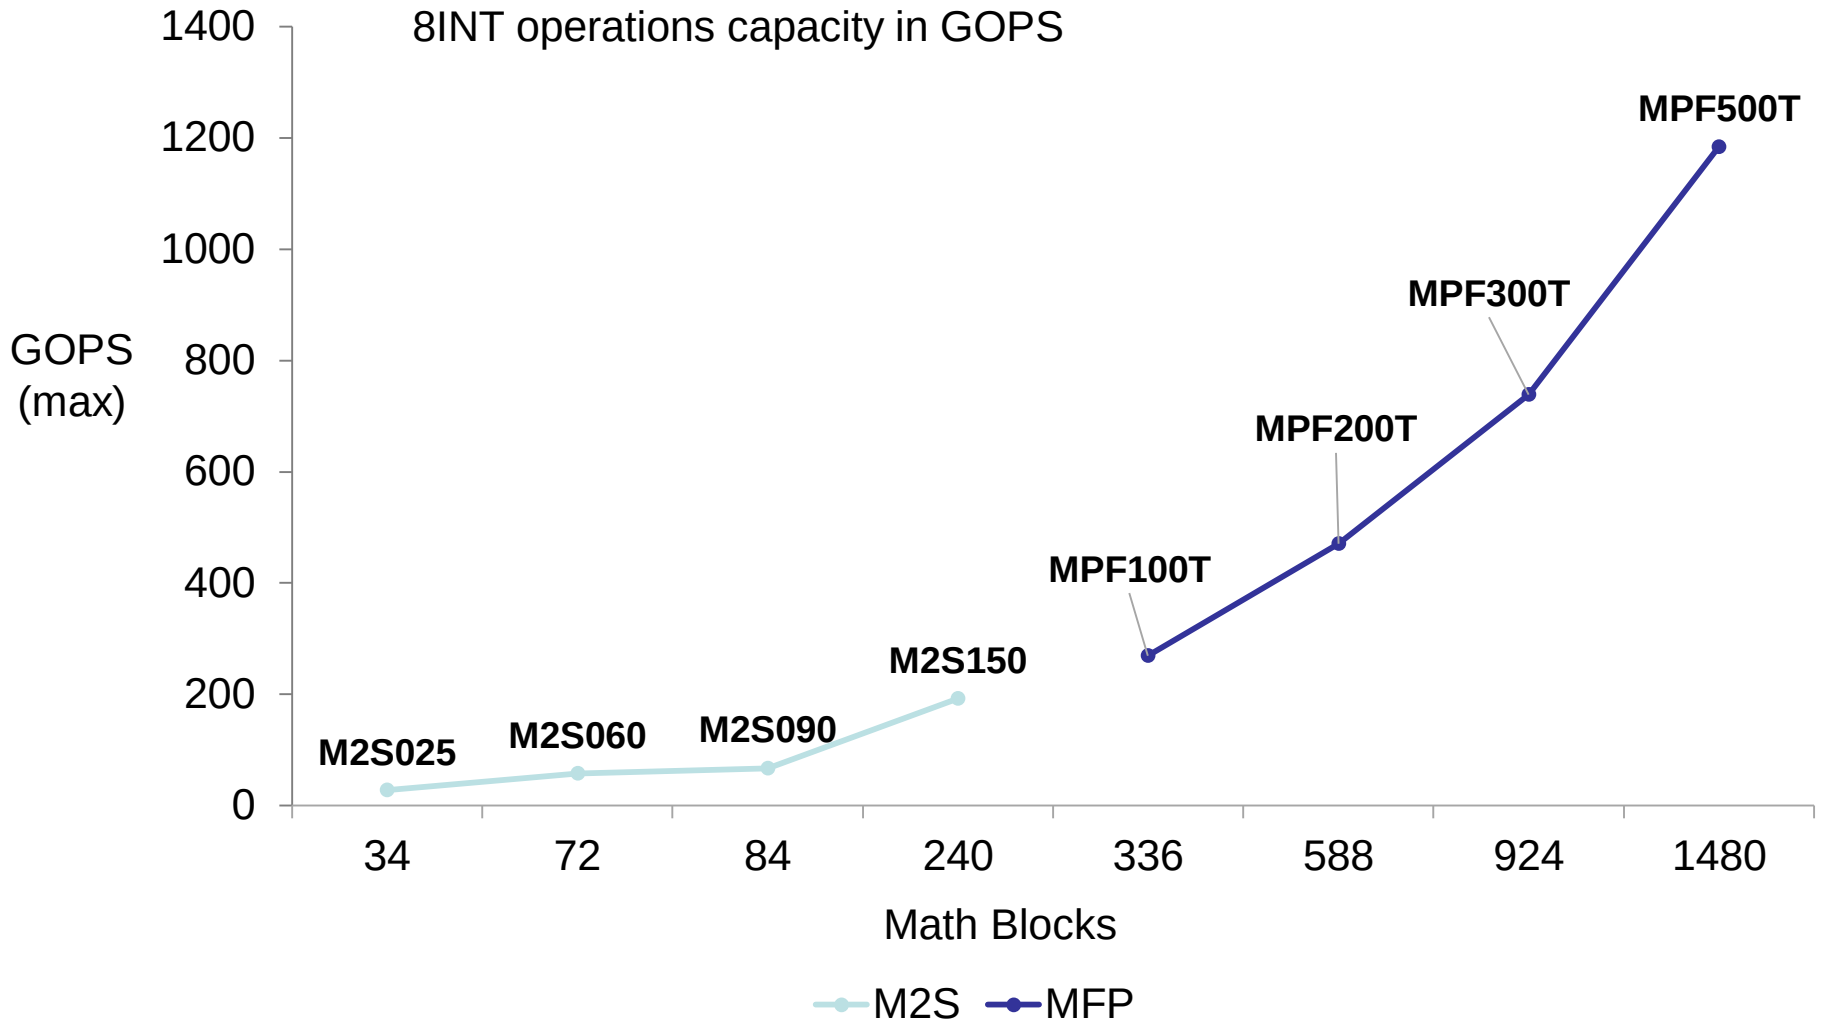

Benchmark Demo

TinyYOLOv2 benchmark demo on MPF300
- Performance 43.6fps at 304 GOPS
- Power consumption 2.98W (Core)
- Resources used ~200K LE

Input image shape	416 x 416
Number of convolutional layers	9
Number of fully connected layers	0
GOPs (MULACC)	7

Runtime (ms)	23.2
Frames Per Second	43.6
Performance [GOPs/s]	304
Power (W)	2.98
Efficiency [GOPs/s/W]	**102**
Multiplier Efficiency [GOPs/s/Slice*]	0.422

SEV Goal – Offer Total System Solution

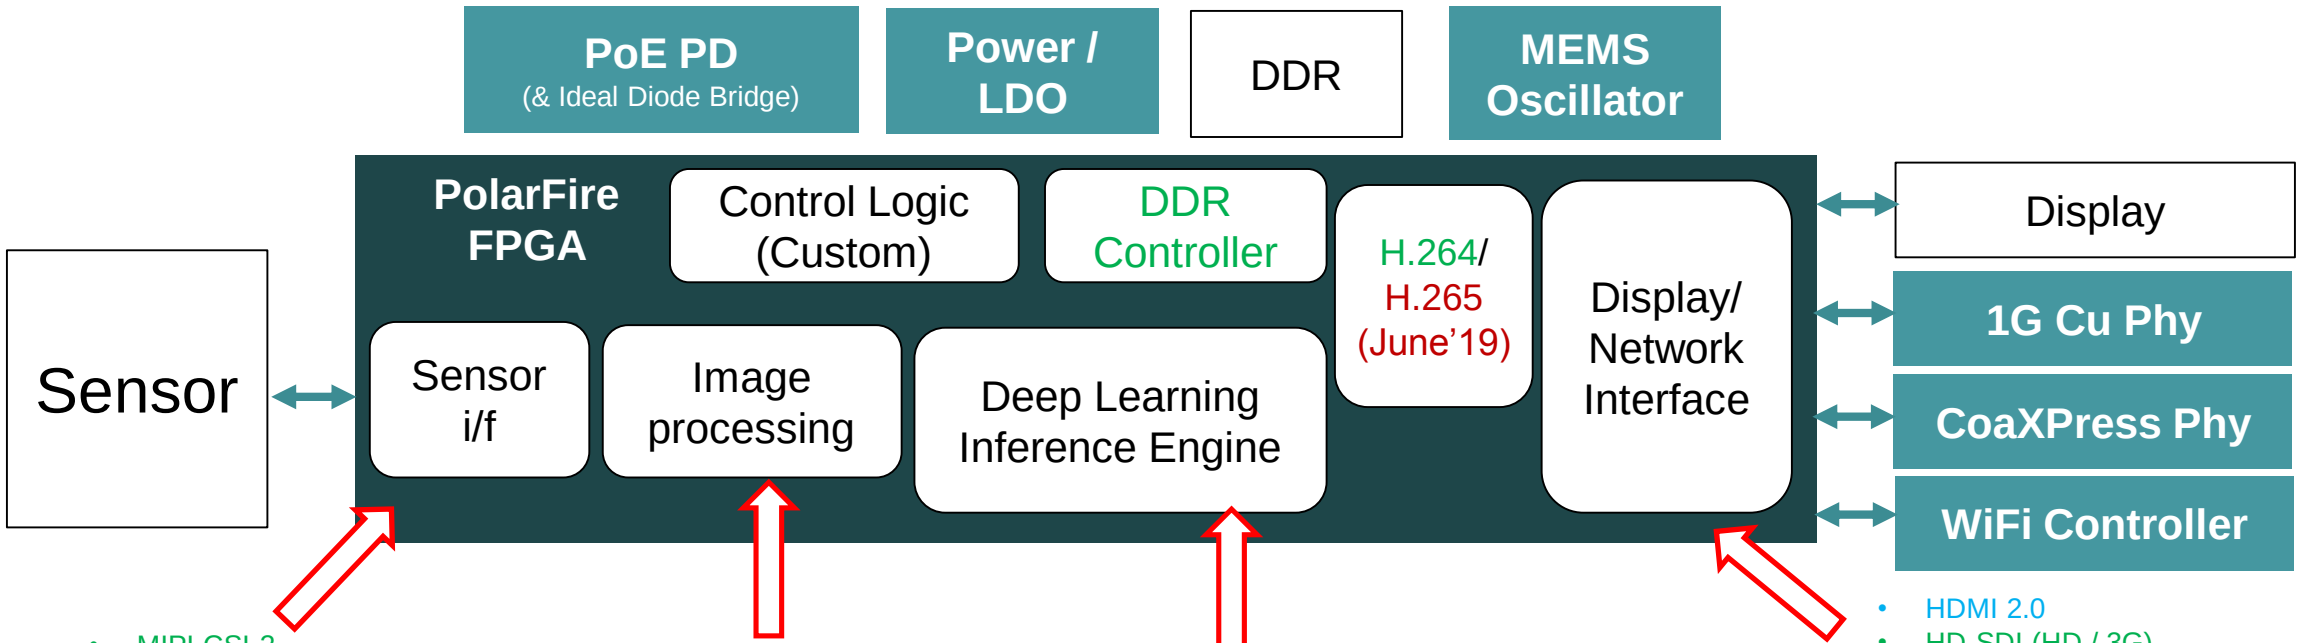

- MIPI CSI-2
- SLVS-EC (July 2019)
- Parallel bus

- Alpha Blending
- Bayer Interpolation
- Display Controller
- Image Edge Detection
- Image Enhancement
- Image Sharpen
- MIPI CSI-2 Rx/Tx
- RGB to YCbCr
- YCbCr to RGB
- Test Pattern Gen

- CNN trained on Caffe
- 2x to 4x performance / watt over competitive FPGA
- TinyYOLOv2 delivering up to 65fps with 416x416 image
- Also support Mobilenet, VGGNet and many others

ASIC Design Services

- HDMI 2.0
- HD-SDI (HD / 3G)
- HD-SDI (6G / 12G) (August 2019)
- CoaXPress (6.125/ 12.5) (June 2019)
- HDCP 2.0
- DisplayPort 1.4a
- USB 2.0 (May 2019)
- USB3.1 gen2 (August 2019)
- LVDS 7:1
- 10G MAC
- 1/2.5/5/10G USXGMII (April 2019)

Available In-house / Available with Third Party / In Development

DEEP LEARNING INFERENCE IMPLEMENTATION

What do we bring to the table

Training of Network (Cloud Based)

Network model (Tiny Yolov2, ResNet50, VGG16, LeNet5…)

Caffe*

Network training

Trained network model

FPGA Implementation

Implementation Constraints for Inference Resources, Power, Footprint

CDL Core Deep Learning — **Network model to RTL compiler**

Libero System-on-Chip

Optimized RTL

Microchip PolarFire™ FPGA

*Caffe is supported now.
*Tensorflow support is planned
*ADS can guide Tensorflow to Caffe conversion

ASIC Design Services

ADS' can provide consultancy

- Feasibility Study
- Specification
- CNN Model and Training

ADS' Core Offering

- CNN Quantization and Fine Tuning
- Core Optimization
- Core Implementation

Engagement Model with ASIC Design Services

To help get started ASIC Design Services provides:

- Help with demo bring up
  - Needs PolarFire Splash Kit + ADS Demo Files (https://www.microsemi.com/existing-parts/parts/144001)
- Webex presentation and training on ADS tools (NDA required)
- Direct customer support (NDA required)

Two Models:

- Single Use – One design per license
- Subscription Model – Multiple designs within a time period

CDL Framework License:	Price USD
Single use	$ 50,000
Subscription 1-year	$ 90,000
Subscription 2-year	$ 150,000

What is Core Deep Learning

- FPGA CNN inference engine – IP Core
- Network specific core
- Platform optimized core
- 8-bit precision
- Comprehensive CNN structure and layer support
- Automated software flow from trained CNN to FPGA design

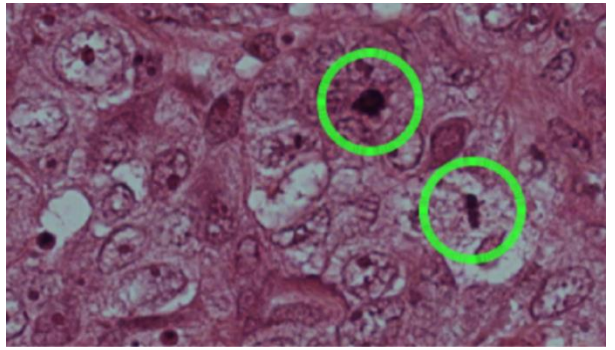

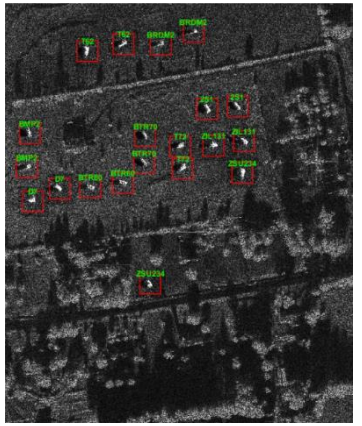

Microsemi
Power Matters.
PolarFire™
FPGA
Cost-optimized
Lowest Power
Mid-range Densities

CDL Flow Overview

	Phase	Description	Processing Platform	Output	Expertise
	Network description	Standard machine learning flow	CPU/GPU	Trained deep neural network	Machine Learning
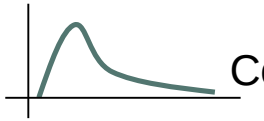	Compression	Proprietary CDL quantization flow – automated	CPU/GPU	Quantized trained CNN – **Bit-accurate** to FPGA Design	Machine Learning
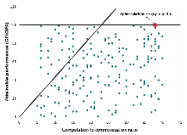	Design space Exploration	Proprietary CDL design flow – automated	CPU	Accurate performance and resource usage report together with SystemVerilog CDL core and memory map.	Software
	CDL SystemVerilog core integration	Microchip FPGA flow	CPU	Bitstream	Hardware

CDL – 8b INT Quantization

Industry Research

Accuracy on a variety of datasets (Gysel, 2016)

Network	Floating-point	8-bit
LeNet	99.1	99.1
CIFAR-10	81.7	81.4
CaffeNet	56.9	56.0

Top5 Accuracy in ImageNet (Guo et al, 2016)

Network	Floating-point	8-bit
CaffeNet	77.12	76.64
VGG16	88.10	87.60
GoogLeNet	88.82	88.64
SqueezeNet	79.72	79.16

ADS Evaluation

Accuracy on a variety of networks/applications (ADS work)

Network	Floating-point	8-bit
LeNet	99.13	99.12
Scene Labelling	73.89	73.31
VGG16	88.44	87.54

Core Deep Learning Status

CNN Layers Supported

- Convolutional layer
- Depthwise separable convolutional layer
- Pooling layer
- Non-Linear activation layer
- Concatenation layer
- Fully-connected layer
- Residual layer

Tested and Validated CNNs

ResNet50	Image Classification
VGG16	Image Classification
SqueezeNet	Image Classification
Cifar10	Image Classification
Tiny YOLOv2	Object Detection
LeNet5	Digit Classification
OpenPose	Pose Estimation
Origami	Scene Labelling
CNNGestureRecognizer	Hand Gesture Recognition

CDL Interfaces

- **Simple APB control interface**
- **Standard AXI memory interface**

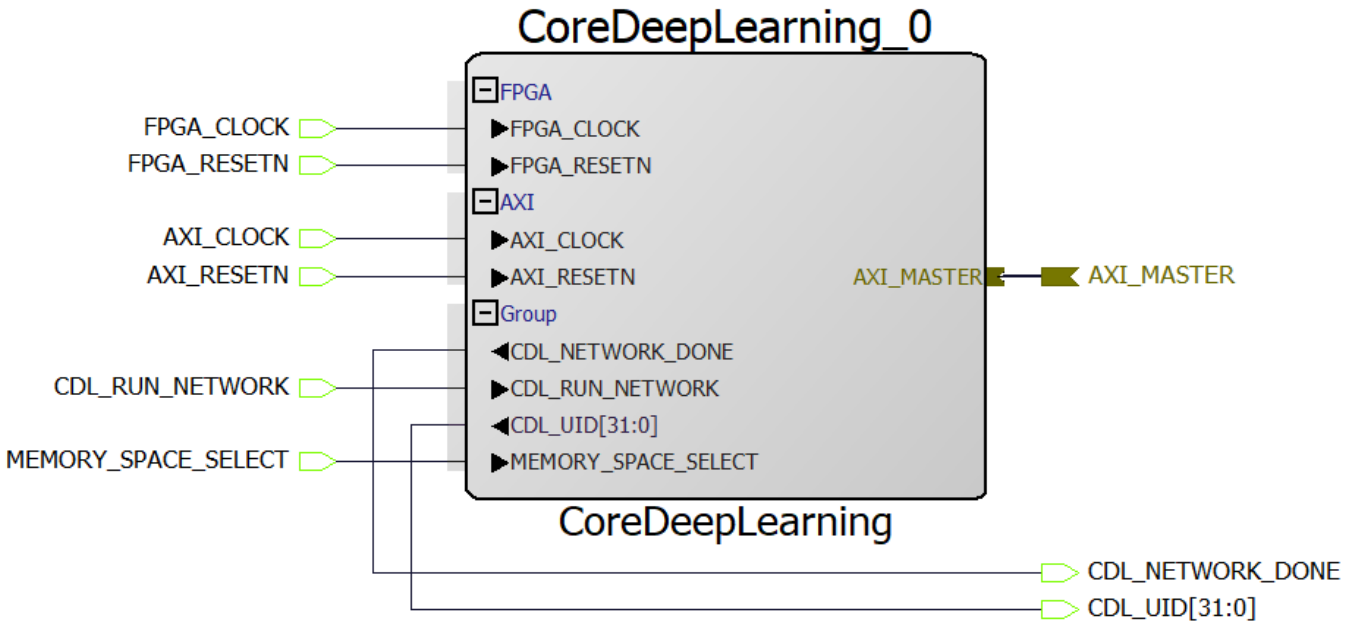

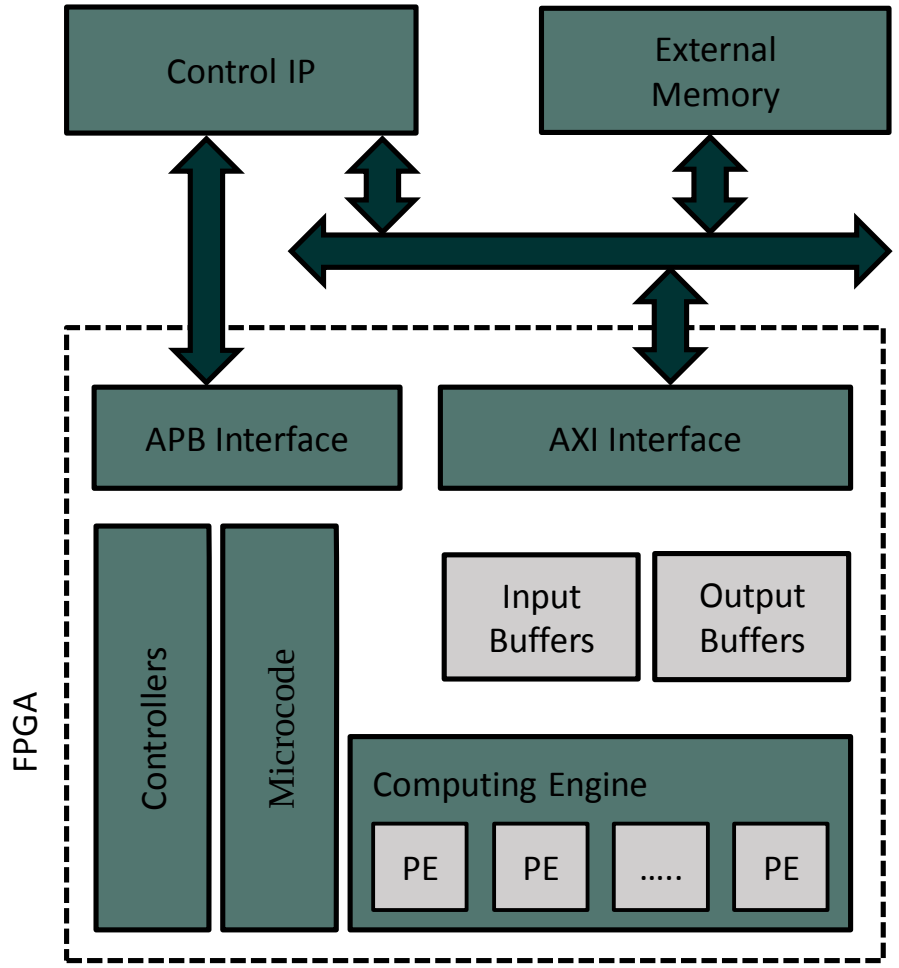

IP LIBRARY, HARDWARE AND KITS

SEV Goal – Offer Total System Solution

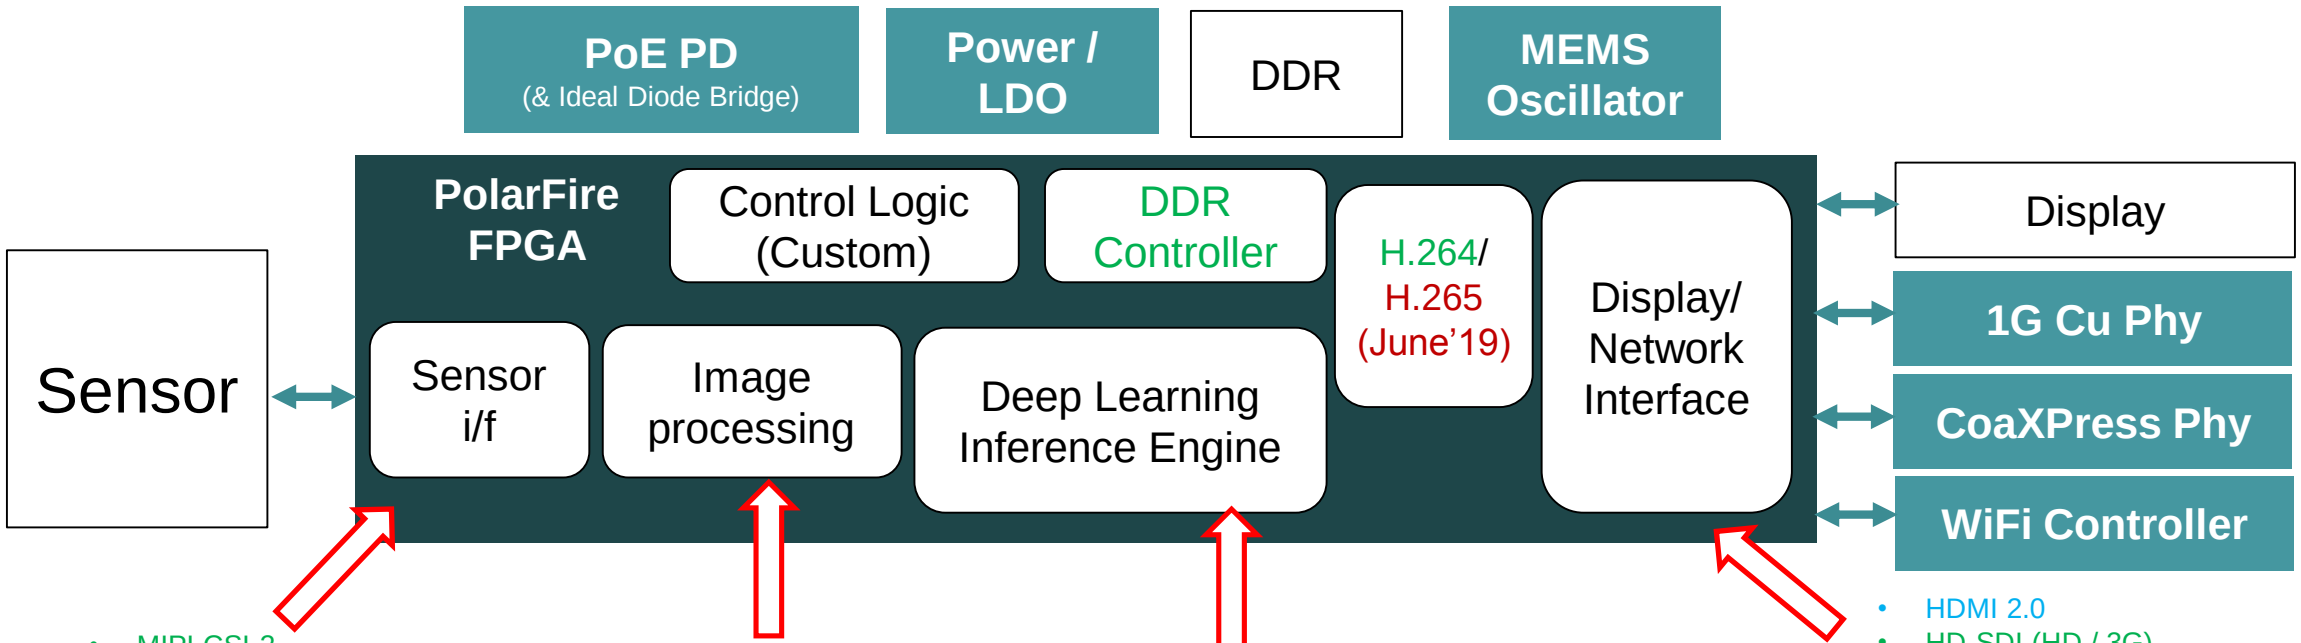

- MIPI CSI-2
- SLVS-EC (July 2019)
- Parallel bus

- Alpha Blending
- Bayer Interpolation
- Display Controller
- Image Edge Detection
- Image Enhancement
- Image Sharpen
- MIPI CSI-2 Rx/Tx
- RGB to YCbCr
- YCbCr to RGB
- Test Pattern Gen

- CNN trained on Caffe
- 2x to 4x performance / watt over competitive FPGA
- TinyYOLOv2 delivering up to 65fps with 416x416 image
- Also support Mobilenet, VGGNet and many others

ASIC Design Services

- HDMI 2.0
- HD-SDI (HD / 3G)
- HD-SDI (6G / 12G) (August 2019)
- CoaXPress (6.125/ 12.5) (June 2019)
- HDCP 2.0
- DisplayPort 1.4a
- USB 2.0 (May 2019)
- USB3.1 gen2 (August 2019)
- LVDS 7:1
- 10G MAC
- 1/2.5/5/10G USXGMII (April 2019)

Available In-house / Available with Third Party / In Development

SEV – Video Platform

PolarFire Video Kit (MPF300-Video-Kit)

- 8MP 4K Sony (IMX334) Dual Camera Interface
- MIPI CSI-2 Tx/ Rx, DSI Connection
- HDMI 2.0 Rx/Tx, HDMI 1.4 Tx
- HPC FMC Connection
- HD-SDI Rx/Tx

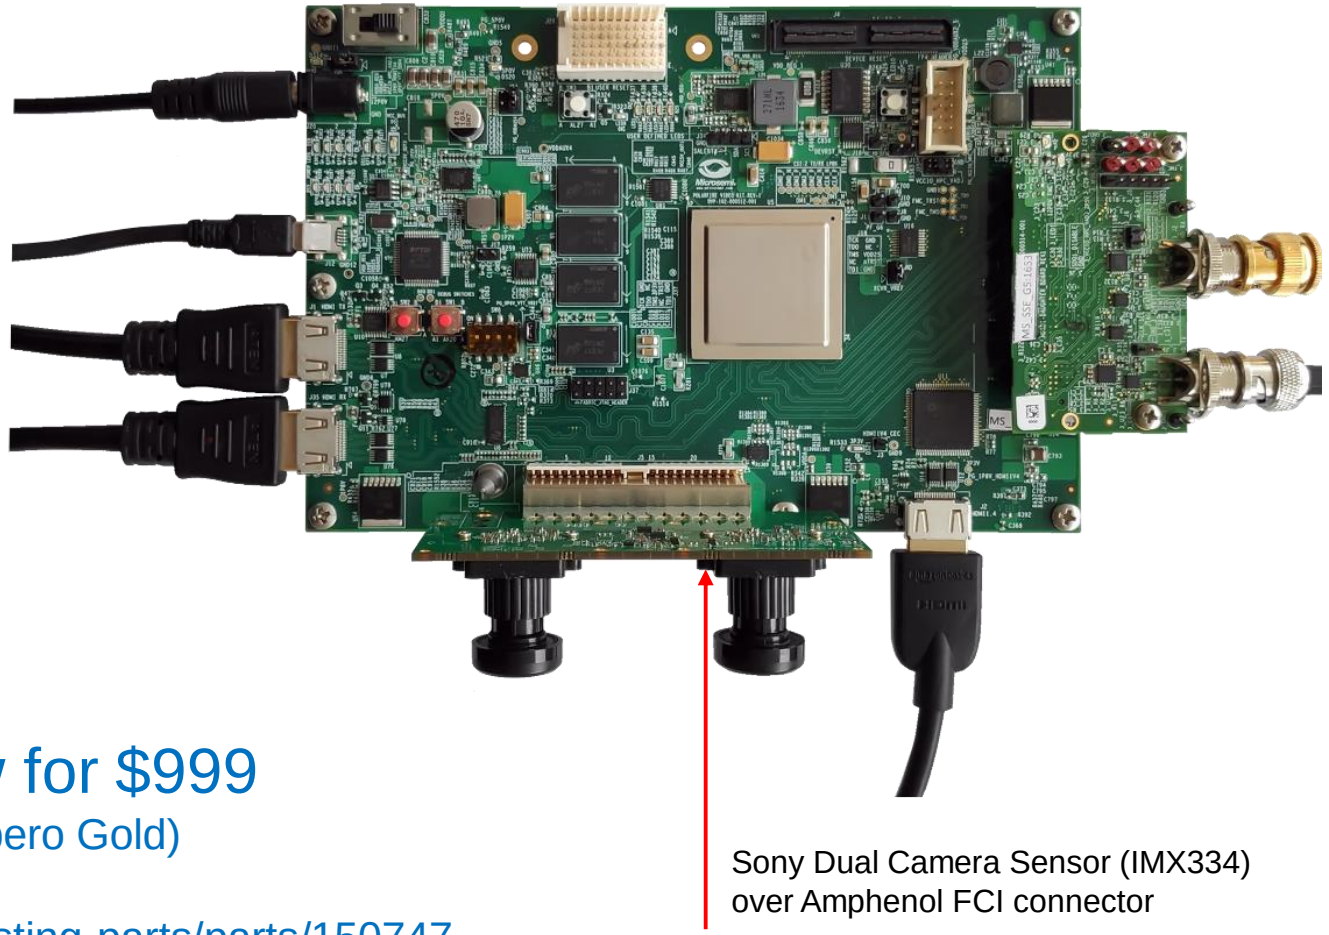

Sony Dual Camera Sensor (IMX334) over Amphenol FCI connector

Other Resources

- Total 2GByte DDR4 SRDAM (4 x 256*16Mbits)
- USB to UART
- 1Gb SPI Flash

Available Now for $999

(includes free Libero Gold)

https://www.microsemi.com/existing-parts/parts/150747

Ideal Platform for Video and Imaging

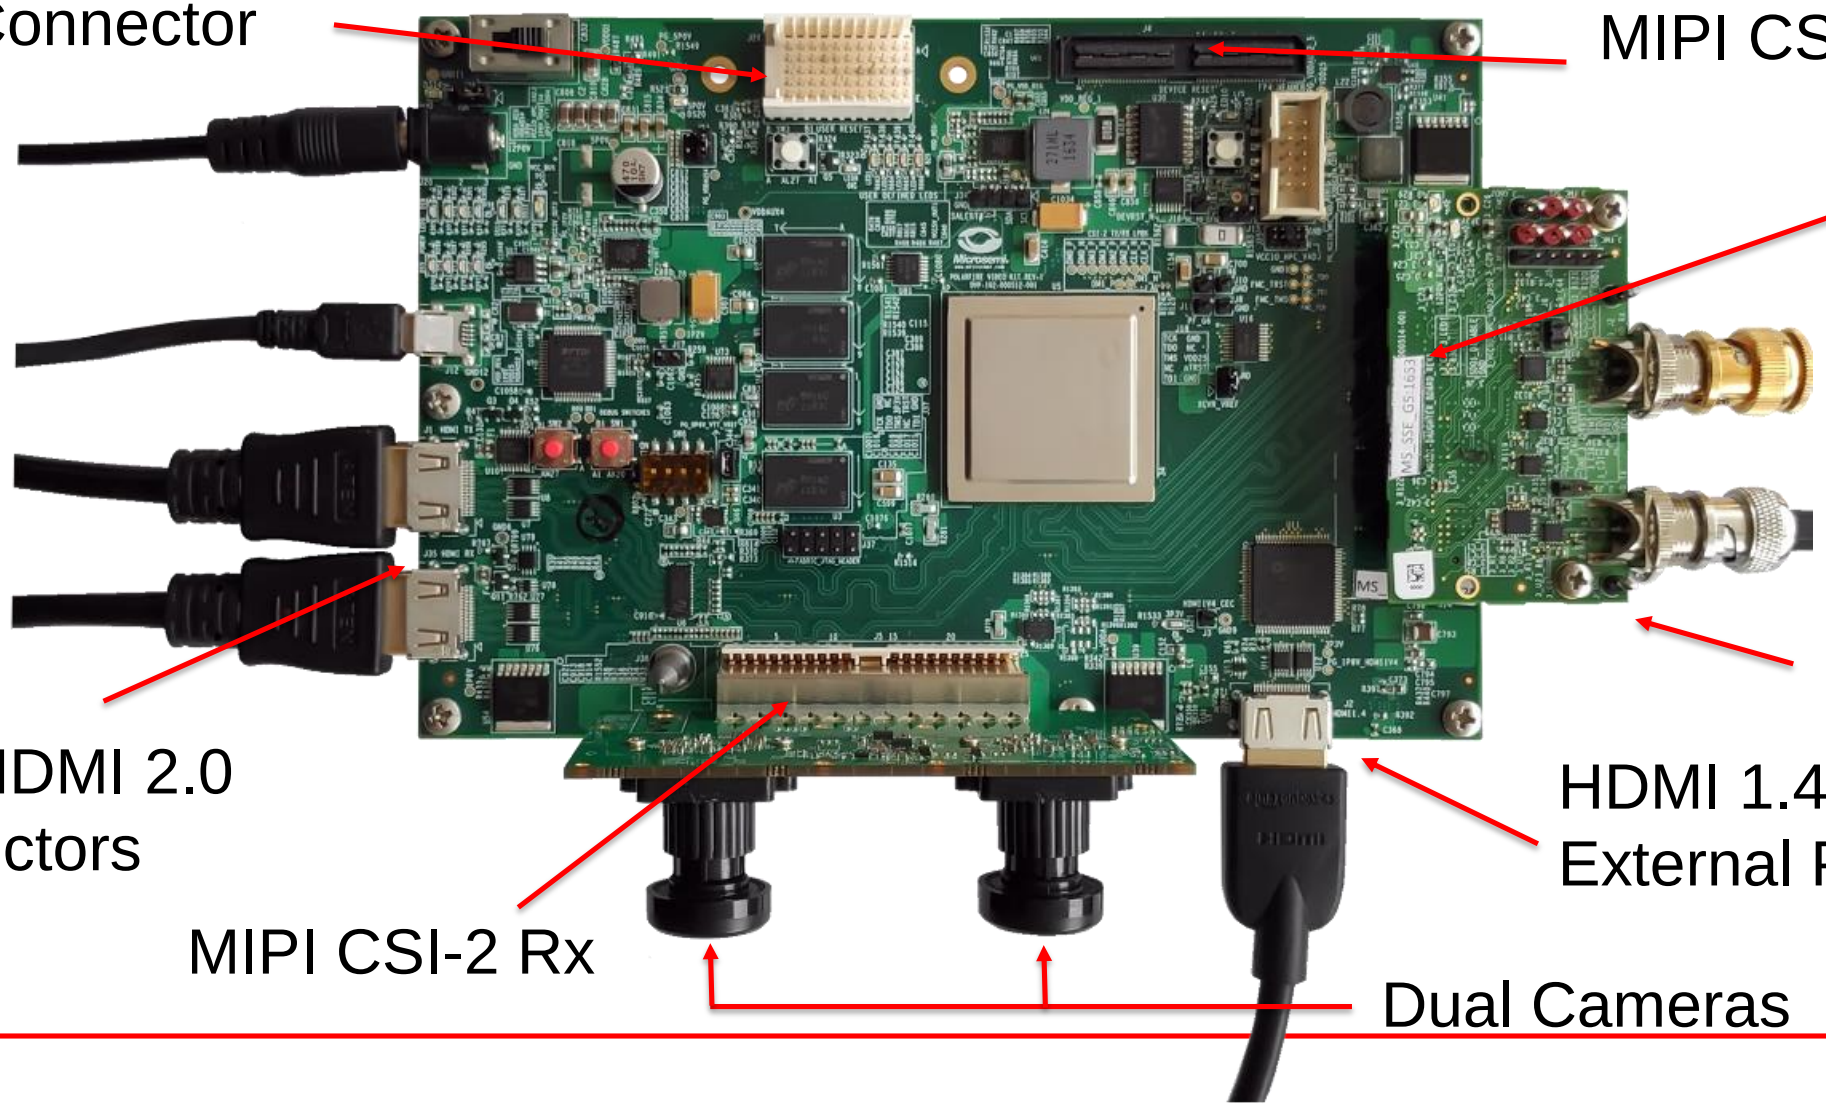

Demonstration

Picture in Picture Demo

← Bayer

Edge Detection →

Picture 1 : 1920 x 1080
Picture 2 : Scaled 1/3

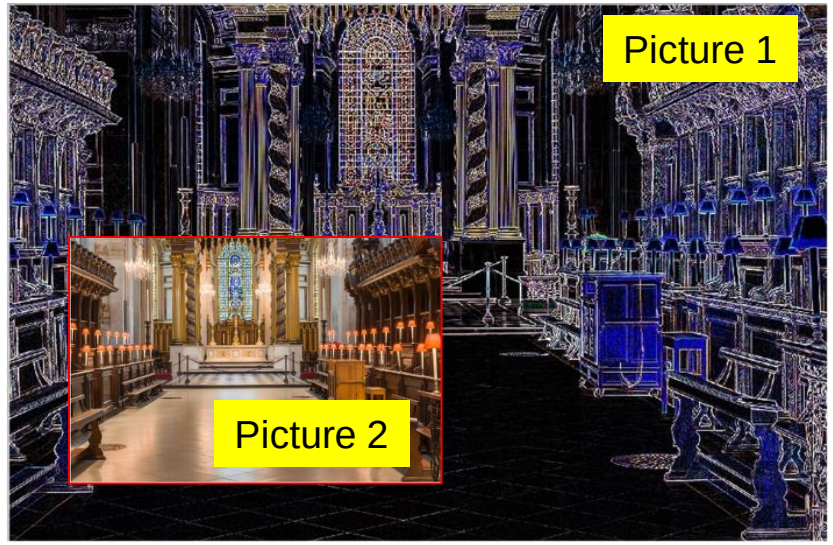

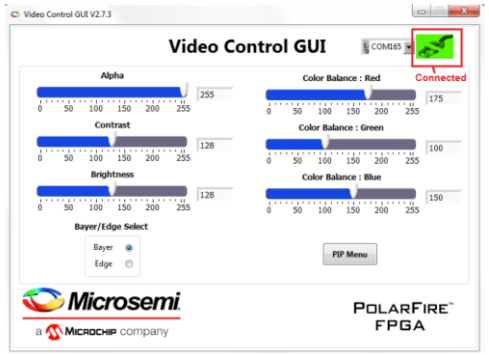

Video Control GUI
Alpha Blending
Contrast
Brightness
Color Balance: RGB

PIP Menu
Select Source
Choose movement speed

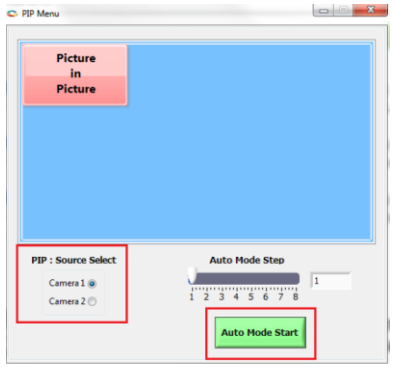

Demonstration Flow

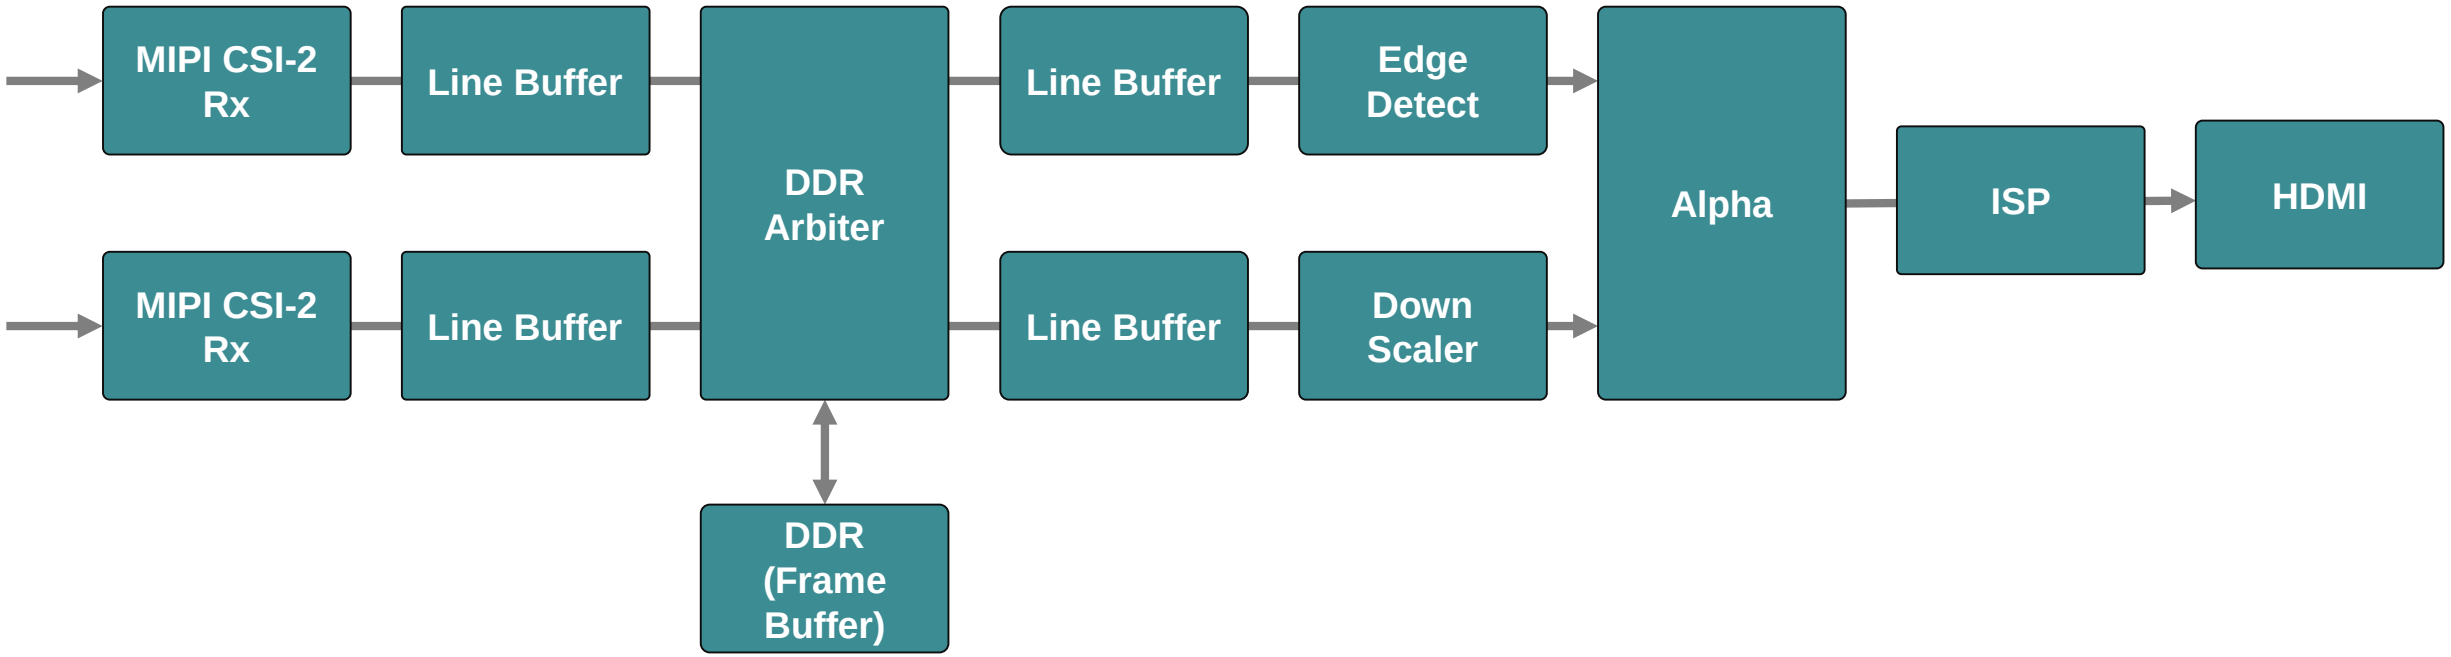

- Displays two 1080p video streams in picture-in-picture format
- Edge detection and down scaling
- Combine video outputs, process the video stream and display via external HDMI PHY

PolarFire Imaging ISP Portfolio

Core Name	Encrypted	RTL	Current Availability
Bayer Conversion	**Free**	**$1499 (Bundle)**	**Available**
Video DMA			
Alpha Blending			
Color Space (YCbCr)			
Color Space (RGB)			
Image Sharpening Filter			
Display Enhancement (Brightness/Contrast/Color balance)			
Image Edge Detection			
Display Controller			
Pattern Generator			
MIPI CSI-2 Receiver Decoder Support for RAW8, 10, 12 format. Tested to 1Gbps / lane (1.2 in char)			
MIPI CSI-2 Tx Support for RAW8, 800Mbps. RAW10 in char			
LVDS 7:1 Display (TX and Rx)			

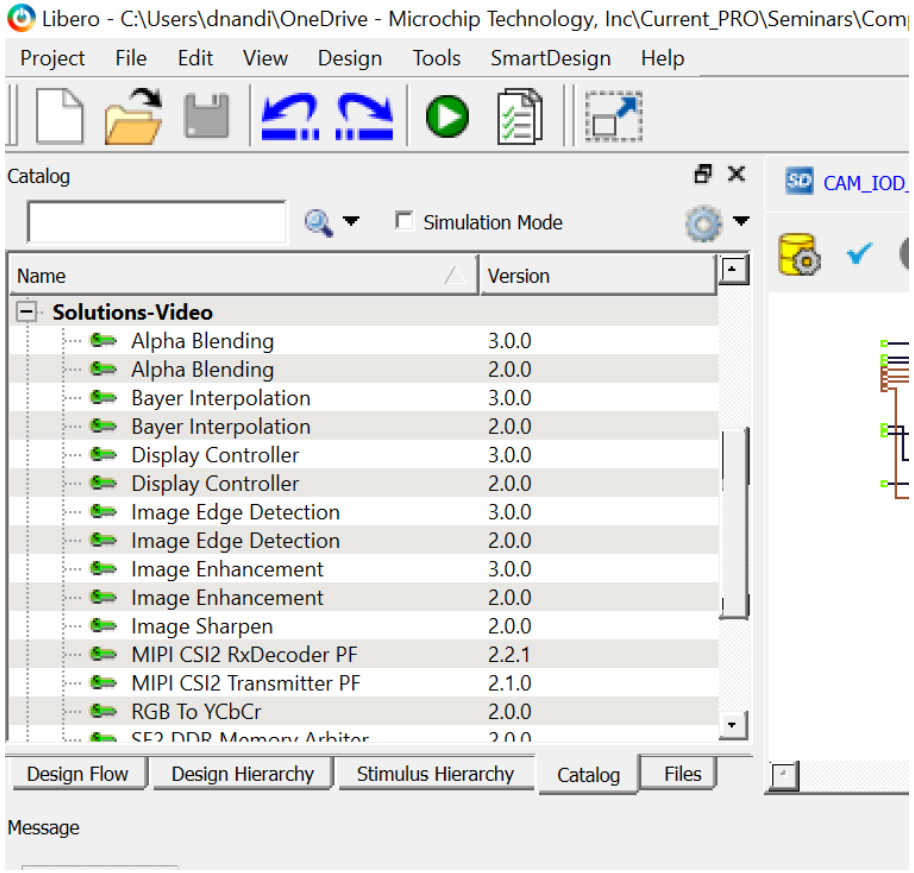

https://www.microsemi.com/existing-parts/parts/150747#resources

- All IPs are part of the Libero Catalog
- Newer versions of IP are pushed via Libero

Microsemi Imaging / Video Kit

- Based on the PolarFire Evaluation Kit (MPF300-EVAL-KIT-ES) Available Now
- Leverages the PolarFire Kit via the Imaging and Video FMC Connector Available Now

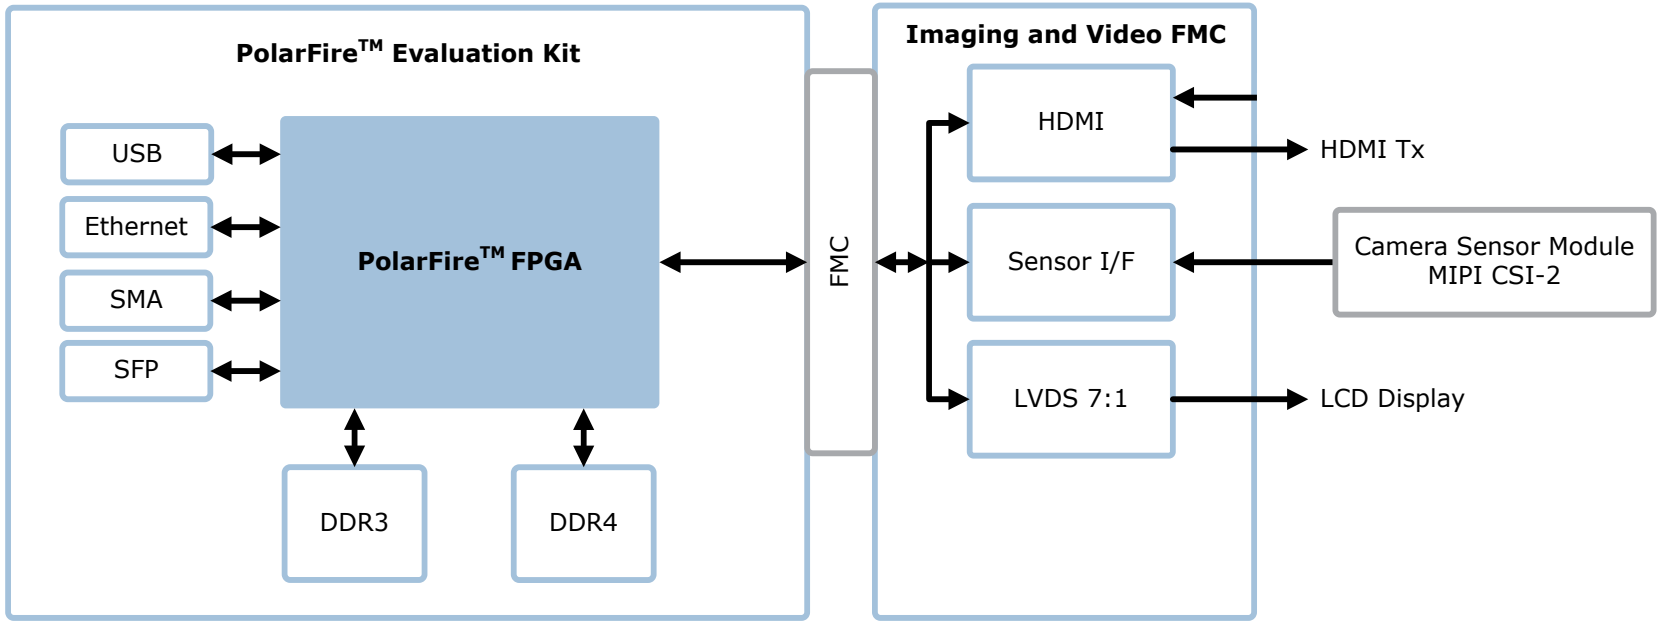

- Also available: Smartfusion2 Advanced Development Kit (SF2 Dev Kit)

Video FMC Boards

- Video FMC with a “MIPI CSI-2 sensor” interface
- Works with PolarFire Evaluation kit and Smartfusion2 Advanced Development Kit

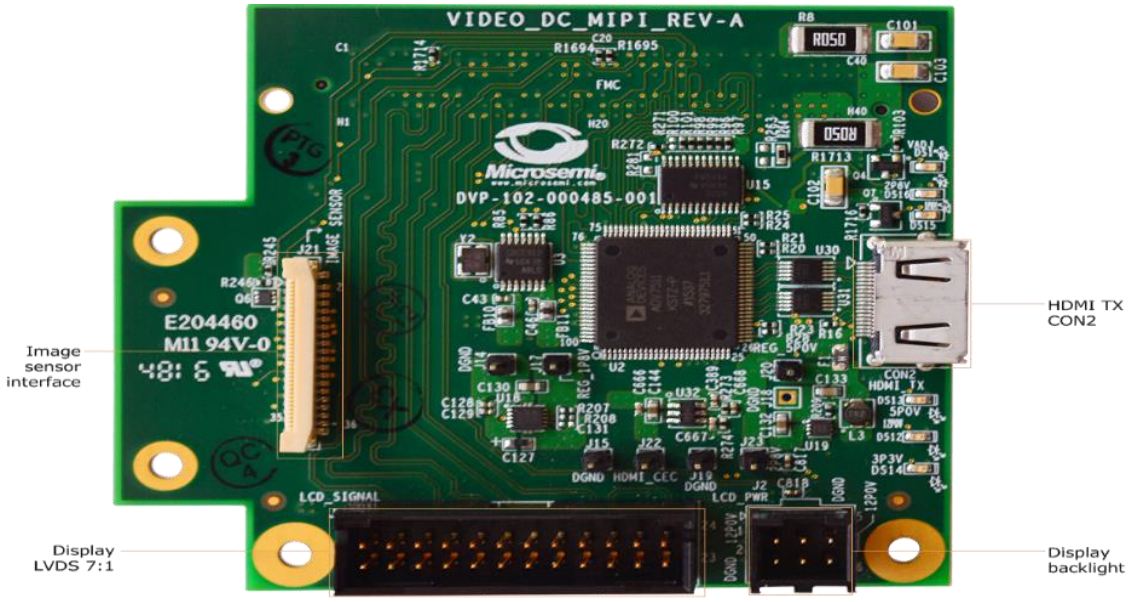

- Video FMC with a “Parallel” interface
- Works with Smartfusion2 Advanced Development Kit

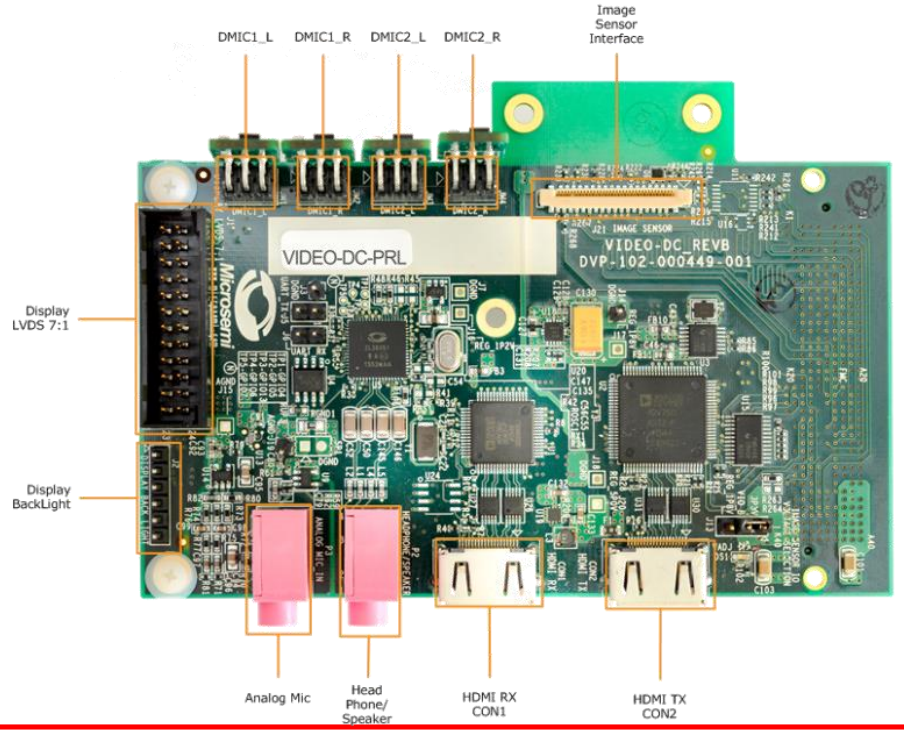

PolarFire/SF2 Imaging/Video Kit

- PolarFire Imaging Kit

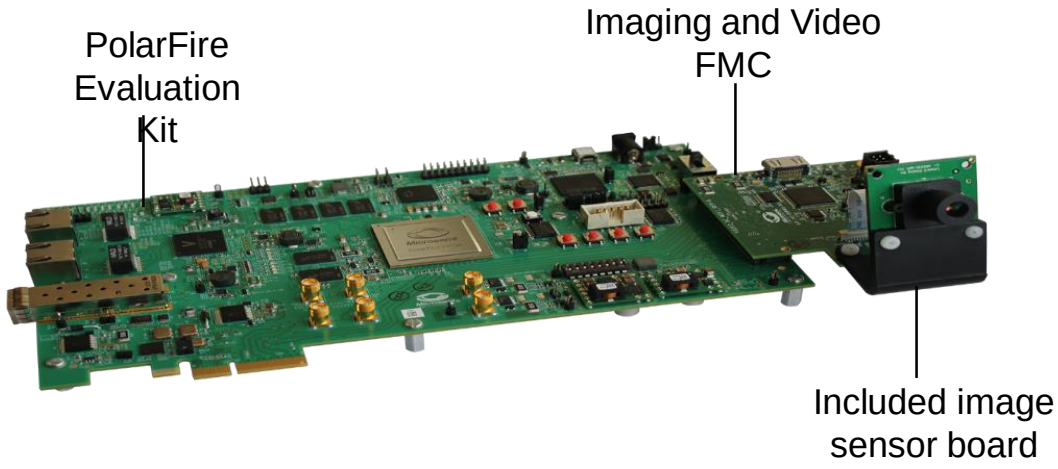

- SmartFusion2 Imaging Kit

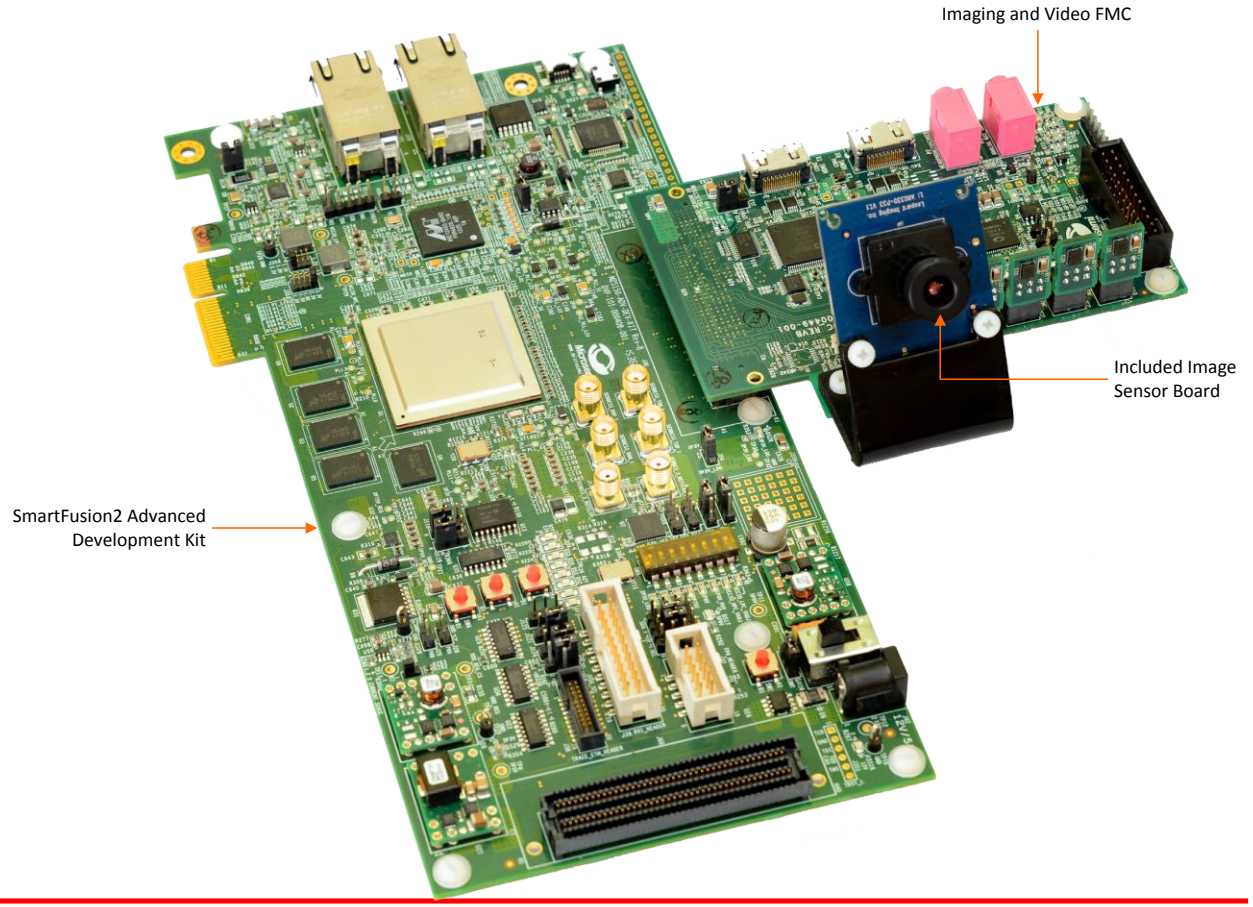

PolarFire SoC Development Platform

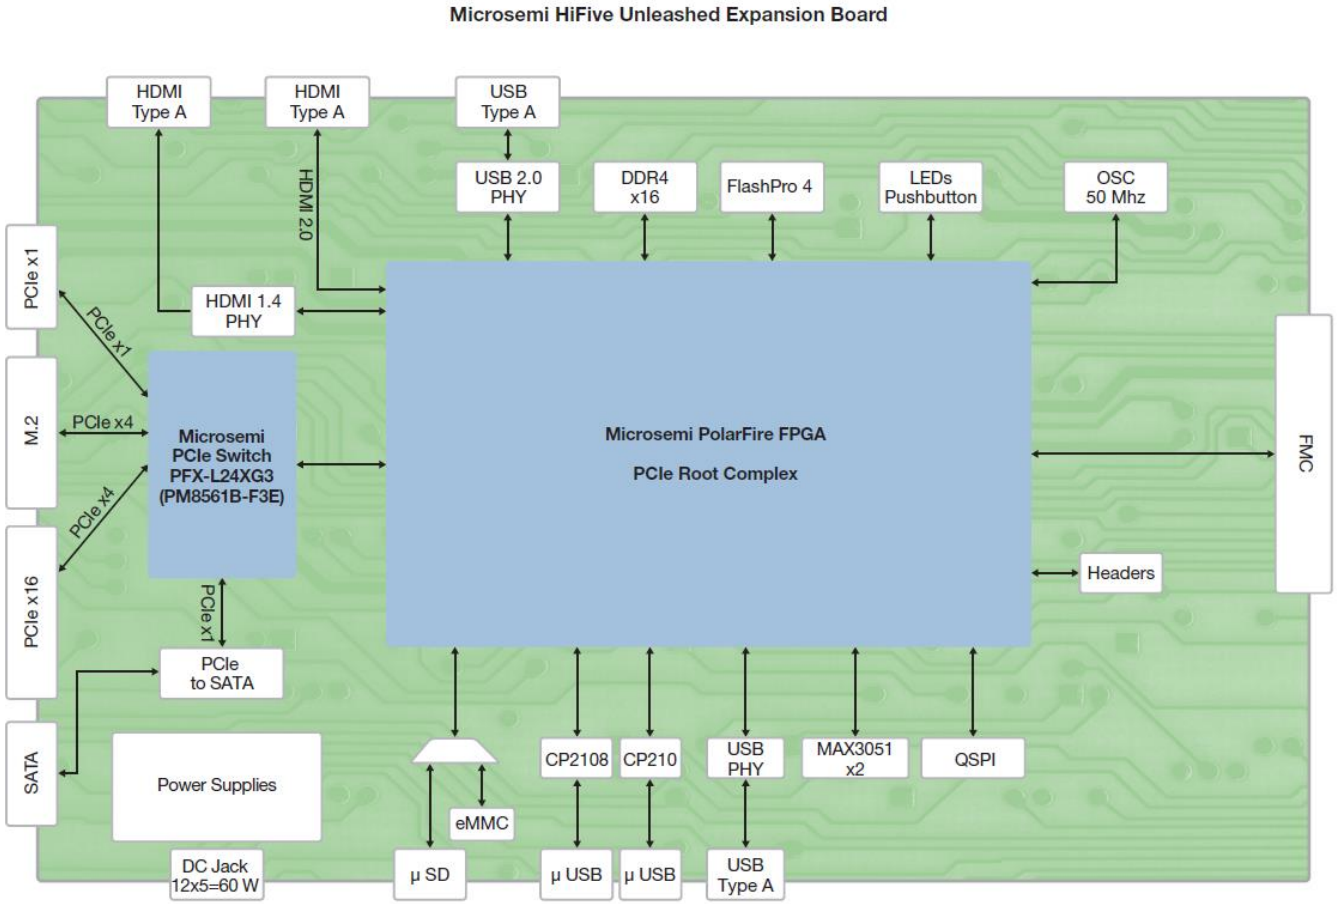

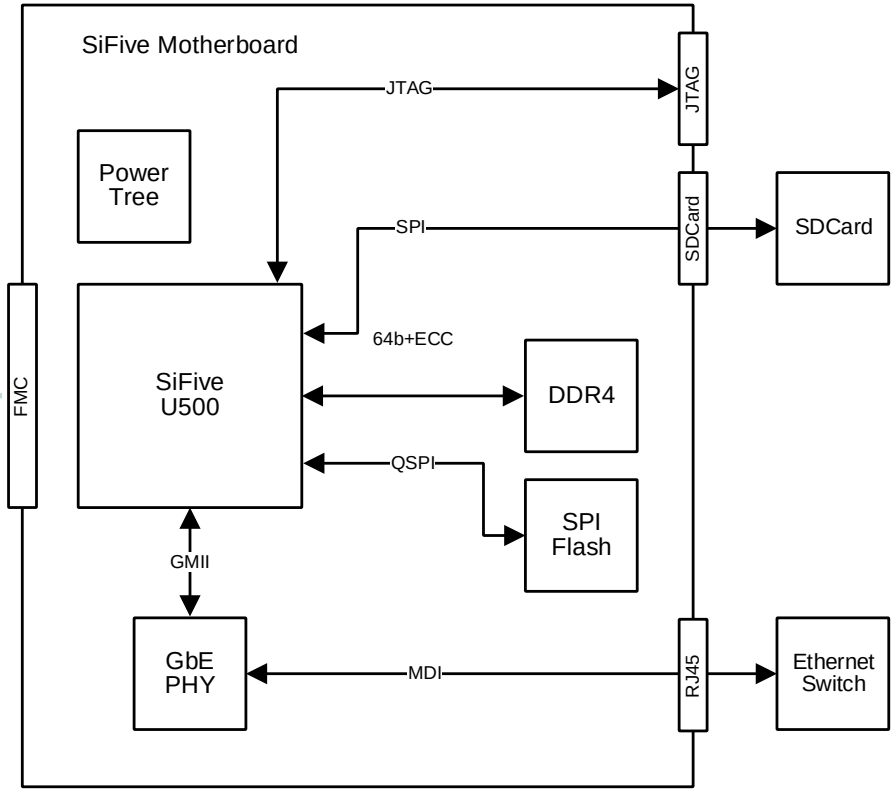

Microsemi HiFive Unleashed Expansion Board

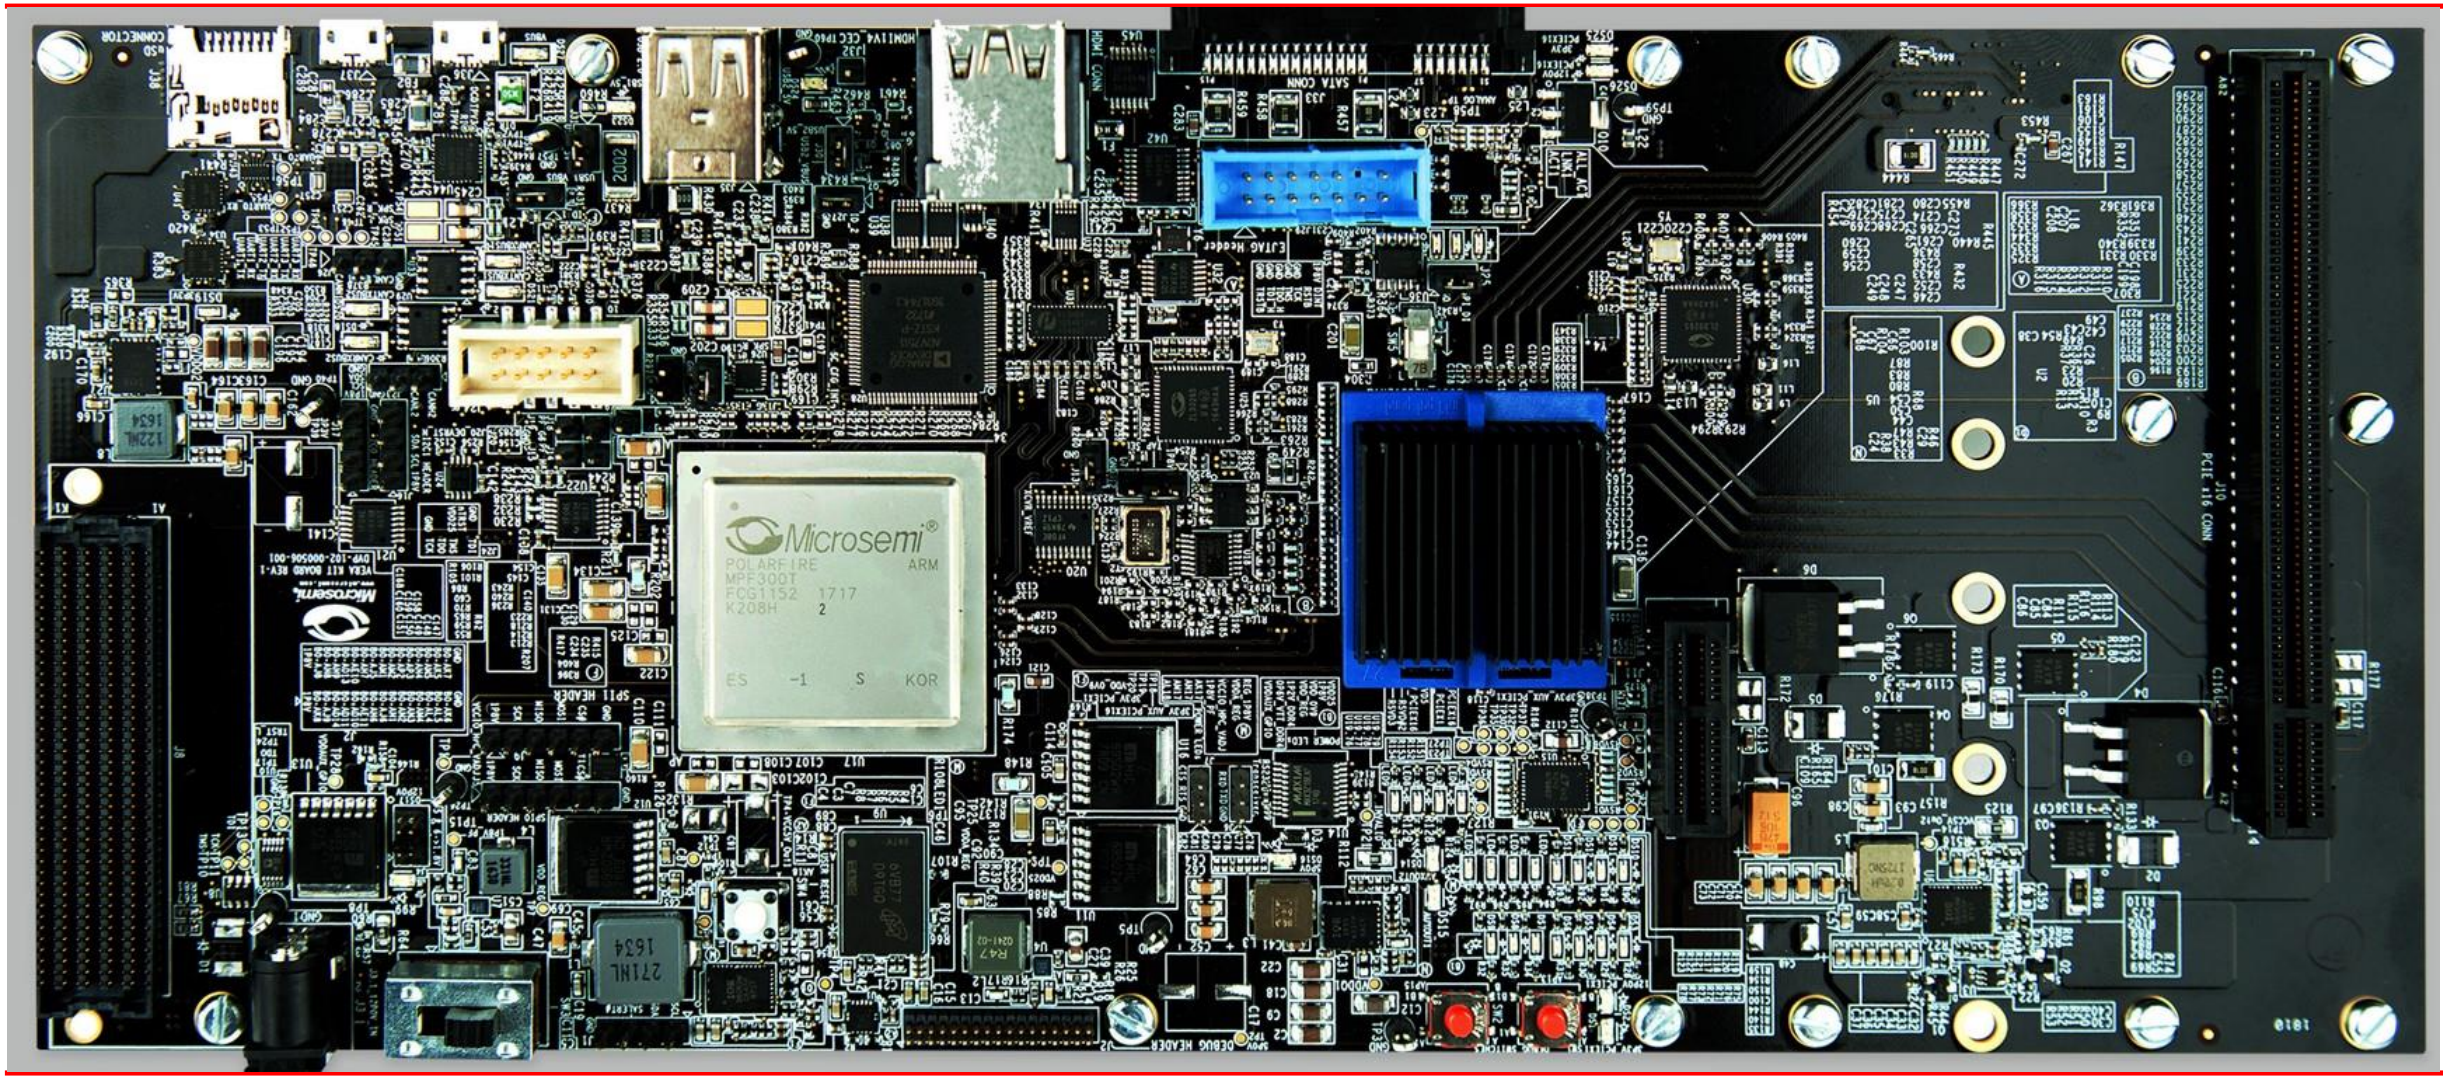

INTRODUCTION TO POLARFIRE

Why Microsemi FPGA

Technology

Flash Based Configuration

- <1/10th static power at 100C
- SEU immune configuration
- On chip NVM, Secure key storage

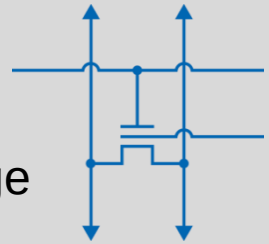

Highly Optimized SERDES

- ~90mW at 10Gb
- Save >120mw per channel

Focused Innovation at Mid Range

- Debug: Smart Debug, Live Probe, Identify Me
- Security: DPA resistant, Authentication
- Embedded: Athena core for Crypto

Benefits

Lowest Total Power Consumption

- Reduce total power by 30% to 50%
- PolarFire: Save 3W to 8W per socket

Lowest Total Cost

- DDR4, 1.6Gbps LVDS
- More resources: IO count, PCIe
- Small form factor packaging

Best In Class Mid Range FPGA

- Productivity: Debugging made easy
- Security: Defense grade security
- Reliability: SEU immune logic for control plane applications

Broad Range FPGA Supplier (100K - 500K LE)

Features	SmartFusion ProASIC3, IGLOO	SmartFusion2 IGLOO2	PolarFire
Logic Elements	100 - 30K	5K - 150K	100K - 480K
Transceiver Rate	--	1-5 Gbps	250 Mbps-12.7 Gbps
I/O Speeds	400 Mbps LVDS	667 Mbps DDR3 750 Mbps LVDS	1600 Mbps DDR4 1.6 Gbps LVDS
DSP (18x18 Multipliers)	--	240	1480
Max RAM	144 Kb	5 Mb	33 Mb
Processor Option	Hard 100 MHz ARM Cortex-M3	Hard 166 MHz ARM Cortex-M3 Soft RISC-V	Soft RISC-V Soft ARM Cortex-M1 Hard Crypto Processor
On-board Flash	Up to 512 KB code store	Up To 512 KB code store	56 KB secure NVM
Family Type	CPLD Replacements Smallest Packages	Low Density FPGAs with more resources & lowest power	Mid-Range Density FPGAs Lowest Power, Cost Optimized

PolarFire FPGA Architecture

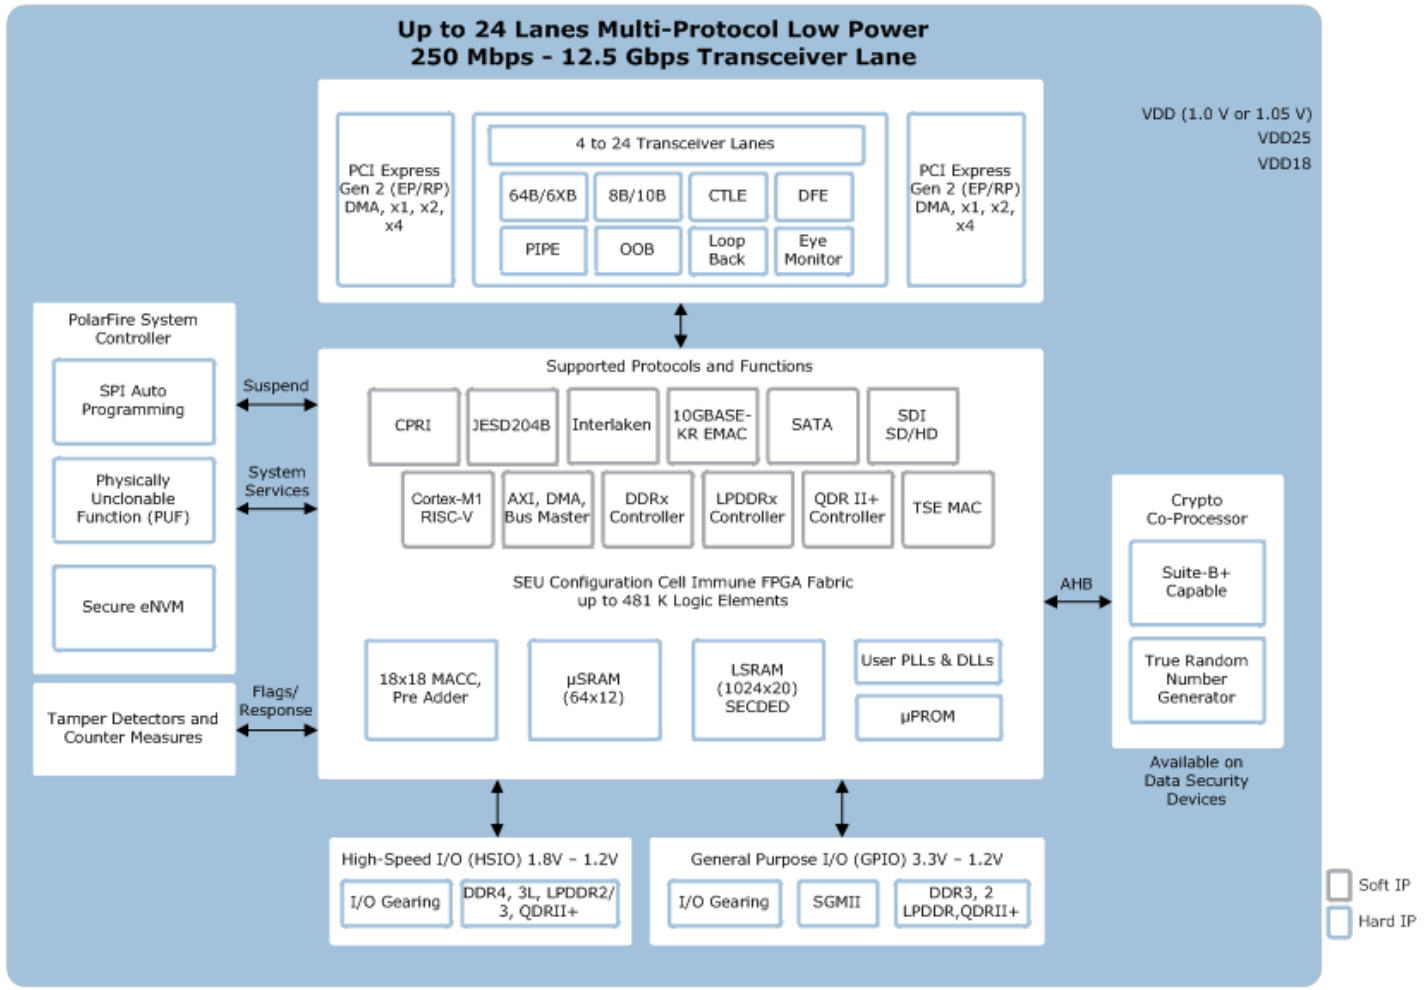

PolarFire Product Family

	Features	PolarFire FPGA			
		MPF100	MPF200	MPF300	MPF500
FPGA Fabric	Logic Elements (4LUT + DFF)	109	192	300	481
	Math Blocks (18x18 MACC)	336	588	924	1480
	LSRAM Blocks (20k bit)	352	616	952	1520
	uSRAM Blocks (64x12)	1008	1764	2772	4440
	Total RAM Mbits	7.6 Mbits	13.3 Mbits	20.6 Mbits	33Mbits
	uPROM Kbits	297 Kbits	297 Kbits	459 Kbits	513 Kbits
	User DLL's/PLL's	8 each	8 each	8 each	8 each
High Speed I/O	250 Mbps -12.7 Gbps Transceiver Lanes	8	16	16	24
	PCIe Gen2 End Points/Root Ports	2	2	2	2
Total I/O	Total User I/O	296	368	512	584
Packaging	**Type / Size / Pitch**		**Total User I/O (HSIO / GPIO) / XCVRs**		
	FCSG325 (11x11, 11x14.5*, 0.5 mm)	170(84/86)/4	170(84/86)/4*		
	FCSG536 (16x16, 0.5 mm)			300(120/180)/4	
	FCVG484 (19x19, 0.8 mm)	284(120/164)/4	284(120/164)/4	284(120/164)/4	
	FCG484 (23x23, 1.0 mm)	244(96/148)/8	244(96/148)/8	244(96/148)/8	
	FCG784 (29x29, 1.0 mm)		368(132/236)/16	388(156/232)/16	388(156/232)/16
	FCG1152 (35x35, 1.0 mm)			512(276/236)/16	584(324/260)/24

Devices in the same package and family type are pin compatible

Extended and Industrial Temperature Support for all Die Package Combinations

PolarFire Architecture Advantage

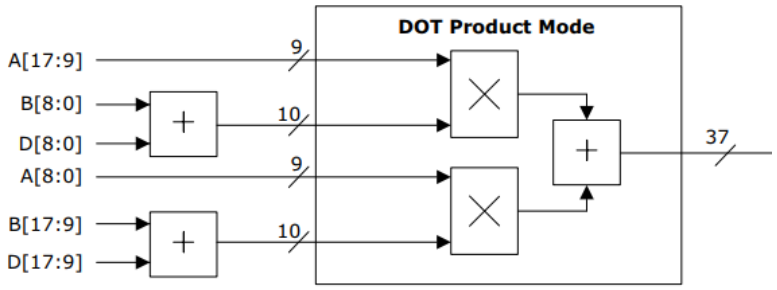

PolarFire Math Block

Family	INT8 Ops per DSP / Math block	With Cascade*
PolarFire	4	4
UltraScale+ (DSP48E2)	2	3.5
Series 7 (DSP48E1)	2	3

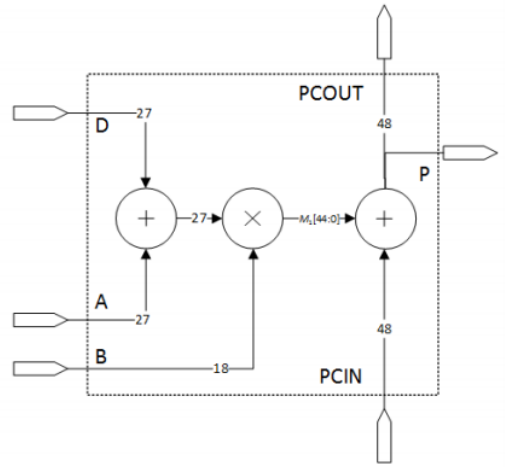

Kintex Math Block

PolarFire: Int8 Operations per DSP / Math block
- 25% Higher vs Kintex7 with Cascade
- 12.5% Higher vs Kintex Ultrascale/+ with Cascade

PolarFire: Power dissipation
- Upto 50% lower total power consumption vs Kintex / UltraScale/+

PolarFire offers Upto
- 2.6 x GOPS **per watt** vs Kintex7
- 2.2 x GOPS **per watt** vs Kintex UltraScale / Kintex UltraScale+

* Cascading needs additional logic and memory; other restrictions

Lowest Total Power – Save Up to 50%

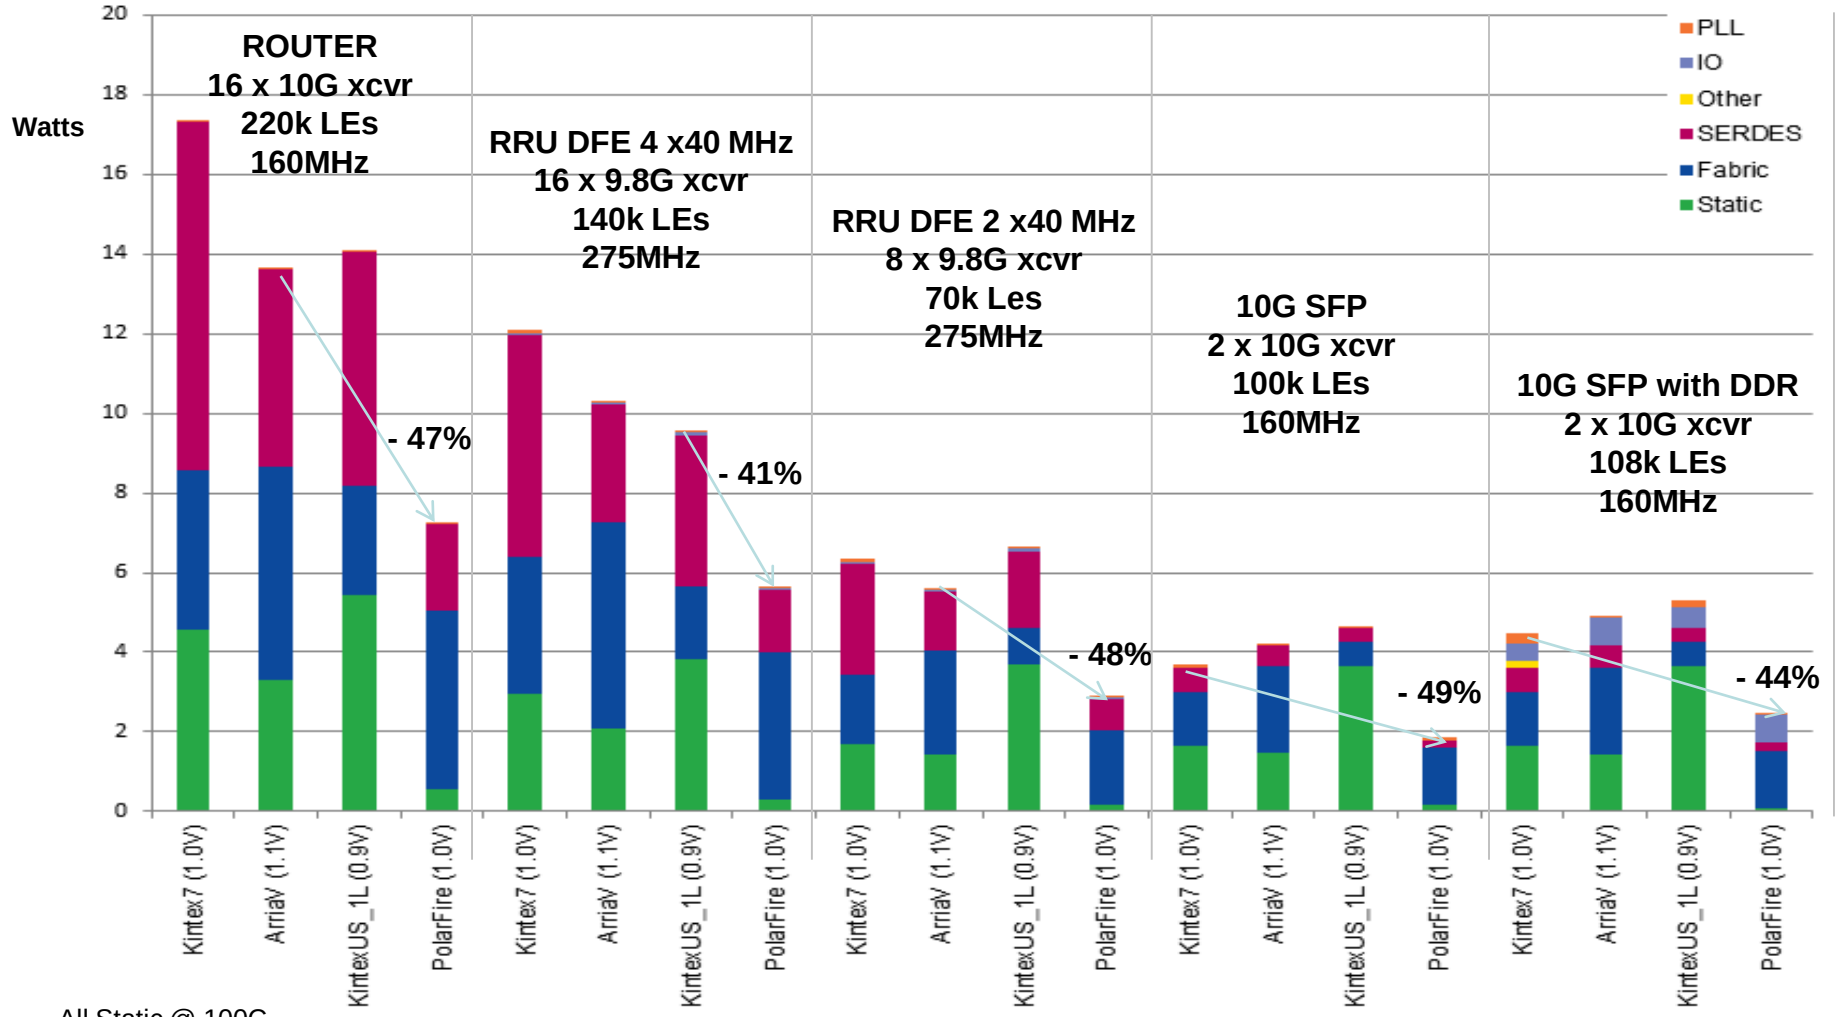

All Static @ 100C

Thank You