


Why FPGA is the Best Platform for Deployment of Machine Learning Application on the Edge from Smart Cities to Space



A Leading Provider of Smart, Connected and Secure Embedded Control Solutions

Puneet Kumar, Manager-Business Development, Aerospace & Defense
Michael Reznik, Staff Embedded Solutions Engineer, Australia

6 June 2023




SMART | CONNECTED | SECURE

Agenda


- **Introduction to PolarFire® FPGAs**
- **Smart Embedded Vision**
- **VectorBlox™ Machine Learning**
- **PolarFire 2**
- **HPSC**

Microchip Products Portfolio



Microchip FPGAs - It's all about FOCUS

Lowest Power, Mid-Range FPGAs and SoC FPGAs



PolarFire® FPGAs

PolarFire Cost-optimized FPGAs Deliver the Lowest Power at Mid-range Densities

- **Cost Optimized**

- Architecture and process optimizations mid-range densities
- Smallest area high-speed I/Os
- Area optimized 12.7G transceivers
- 1.6 Gbps LVDS I/Os supporting DDR4/DDR3/LPDDR3
- Plus hardened I/O gearing logic with CDR (supports SGMII/GbE on GPIOs)
- Lower BOM costs
 - 0°C-100°C at commercial temp prices


- **Lowest Power**

- Lowest static power
- Up to 50% lower power than competing FPGAs

- **Proven Security with Exceptional Reliability**




PolarFire® FPGA Architecture



PolarFire FPGAs cost-optimized, lowest power at mid-range densities

Enjoy Absolute Freedom to Innovate...

PolarFire® SoC



Asymmetric processing with 2X power efficiency


Quad-core 64-bit RISC-V processors

Linux® and real-time operating systems


Highly flexible 2MB L2 cache

With the most comprehensive RISC-V development ecosystem

PolarFire® SoC Block Diagram



PolarFire® SoC has a 6K CoreMark head start!



RISC-V Ecosystem


- Free and open ISA
 - Clean Slate Design
 - Simple, Stable
 - Modular, Extendable

- RISC-V extends Moore's Law
 - Provides a free “architectural” license enabling innovation
 - Customers now can influence the micro-architectural design



Microchip's Mi-V Ecosystem

A Leading Provider of RISC-V Based Solutions



*Microchip's Comprehensive
Ecosystem
to Support RISC-V
Development*

PolarFire® FPGA in Embedded Vision

Driving your NexGen platform

Accelerate Smart Embedded Vision at the Edge Supporting 4K Resolutions with Low-Power 12.7 Gbps SERDES



Virtual Reality



Drive Assist Systems



Drone Vision



Augmented Reality



Medical Vision



Machine Vision



Satellite Vision



Broadcast



Physical Security



Thermal Vision

- **Target next-gen embedded vision customer platforms**
- **Low power, small form factor and portable designs**
 - More thermal head room on the edge
 - Higher resolution – multiple 4K channels
 - Integrate Artificial Intelligence and Machine Learning
- **Alleviating customer concerns on security and reliability on the edge.**


Benefits of PolarFire® FPGAs in Embedded Vision

- **Lowest total power consumption**
 - Reduce total power by 30% to 50%
 - Save 3W to 8W per socket
- **Lowest total cost**
 - DDR4, LPDDR3, 12.7G SERDES, 1.6 Gbps LVDS
 - Small form factor – 11 x 11 mm packages
- **Best-in-class security and reliability in the mid-range (up to 500K LUTs)**
 - No overbuilding or cloning, secure boot
 - SEU immune zero FIT flash configuration

Longer battery life



Smaller heat sink



Eliminate fans



Smallest form-factor, high-end features, in a mid-range FPGA




MPF200
11 x 14.5mm



MPF300
16 x 16mm

99.999% Reliability




Best-in-class security™





The Licensed DPA Logo is a trademark of Rambus Cryptography Research, Inc., used under license.

Proven Low-power Advantage



- Save \$10 in BOM for fans
- Fan less small form factor 4K Cameras
- No fan noise in professional cameras with mic
- Lowest thermal noise in Infra-red imaging

More Thermal Head Room



Power Estimator Analysis

16 Channel AVB Switcher Power Estimator Results


AVB Switcher (4 Ch.)	Static (W)	Core Dynamic (W)	I/O* (W)	Transceiver* (W)	Total Power (W)
Polar Fire (MPF500TS)	1.493	2.474	2.991	0.769	7.727
Polar Fire (MPF500TSL**)	1.059	2.474	2.991	0.769	7.293
Kintex 7(XC7k325T-2)	1.796	5.471	3.971	4.079	15.317
Arria V(5AGXFB3H)	2.572	4.350	5.451	1.651	14.025
Zynq US+(XCZU7EV-1L(0.72v))	2.864	2.165	4.011	2.694	11.7

*I/O and Transceiver power estimates include static and dynamic components. **Power Optimized Device

16 Channel AVB Switcher Resource Utilization

Resources	Count	Interface	Lanes	I/Os	Rate
4LUT	271K	Transceivers	16	-	2.97Gb/s
DFF	188K				
LSRAM	816				
uSRAM	249	3x72 bit DDR3	-	351	800 Mb/s
MathBlock	500				

Microchip Video IP blocks



• Sensor Interfaces (Rx)

- MIPI CSI-2 Receive (1.5 Gbps / lane)
- SLVS-EC (4.7 Gbps / lane)

• Display Interfaces

- MIPI Tx (1.0 Gbps/ lane)
- MIPI Tx on SERDES (2.5 Gbps/ lane)
- MIPI DSI Tx (1080p60)
- MIPI DSI Tx (40k60)

• Image Processing

- Alpha Blending
- Bayer Interpolation
- Sobel
- Uncanny Edge Detector
- Display Controller
- Image Edge Detection
- Image Enhancement
- Image Sharpen
- RGB to YCbCr
- YCbCr to RGB
- Test Pattern Gen
- Gamma Correction
- Histogram
- White Balance

• Deep Learning Inference

- VectorBlox™ Accelerator SDK v1.4

• Encode (Decode)

- H.264 Partner IP
- mJPEG compression (PolarFire SoC)
- H.264 Baseline/ iFrame (PolarFire SoC)

• DDR Controllers

- DDR4/ DDR3/ LPDDR3
- LPDDR4 Soft Controller

• Transport Interfaces


- CoaXPress 6.25 / 12.5G
- SDI (HD/3G/12G)
- 10G MAC / 10G PHY
- USXGMII (1/2.5/5/10G)
- HDMI 2.0 Rx / Tx
- DisplayPort 1.4 Tx
- USXGMII RSTP demo
- MIPI Tx on SERDES (2.5 Gbps/ lane)
- HDMI 2.0 / 1.4 with HDCP 2.0
- DisplayPort 1.4a with HDCP 2.0
- USB 2.0, USB3.1 Gen1 & Gen2

Available In-house /
Available with Third Party /
In Development In-house (Plan)

VectorBlox™ Accelerator SDK

Software overlay based deployment of ML solutions in PolarFire® FPGAs

The Deep Learning Construct



Network Training


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

Inference

- Need low latency
- Power & space constrained
- Requires security and reliability





FPGA Design Flow



VectorBlox™ Accelerator Software Development Kit


- Enables developers to code in C/C++ and use power efficient neural networks
 - No prior FPGA design experience required
 - Works on Linux and Windows
- Executes models in TensorFlow and ONNX
 - Offers the widest framework interoperability
- Includes a bit-accurate simulator
 - Users can validate the accuracy of the hardware while in the software environment
- Pre-trained Neural Networks demos included
 - Users can load different network models at run time on supported hardware




Available Today




VectorBlox™ Accelerator SDK




Model Optimization and Calibration



VectorBlox™ Accelerator SDK



Using the Simulator




Running Simulation...

265 toy poodle	0.6151123
266 miniature poodle	0.11071777
155 Shih-Tzu	0.07495117

bandwidth per run = 32002180
estimated 0.0400075 seconds at 100MHz




VectorBlox™ Accelerator SDK



CNNs are supported with Tutorials


Tutorial Name	Source Framework	Task	Accuracy Metric	Accuracy Score	V1000 kcycles
<u>mobilenet-v1-1.0-224</u>	caffe	Classification	topk	65.4	1623
<u>mobilenet-v2</u>	caffe	Classification	topk	70	2037
<u>mobilenet-v1-1.0-224-tf</u>	tensorflow	Classification	topk	68.8	1613
<u>mobilenet-v2-1.0-224</u>	tensorflow	Classification	topk	69	1710
<u>mobilenet-v2-1.4-224</u>	tensorflow	Classification	topk	74.8	2625
<u>resnet-50-tf</u>	tensorflow	Classification	topk	75.2	8786
<u>squeezenet1.0</u>	caffe	Classification	topk	58.2	2326
<u>squeezenet1.1</u>	caffe	Classification	topk	58.4	1301
<u>Sphereface</u>	caffe	face_compare			3784
<u>onnx_resnet18-v1</u>	onnx	Classification	topk	71.4	3304
<u>mnist</u>	onnx	Classification			21
<u>torchvision_resnet50</u>	pytorch	Classification	topk	75.2	8909
<u>torchvision_wide_resnet50_2</u>	pytorch	Classification	topk	76.2	
<u>yolov2-tiny-voc</u>	darknet	Object Detection	VOCmAP	54.06	8907
<u>yolov2-tiny</u>	darknet	Object Detection	COCOmAP	27.02	9056
<u>yolov3-tiny</u>	darknet	Object Detection	COCOmAP	39.71	10147
<u>yolov2-voc</u>	darknet	Object Detection	VOCmAP	69.99	
<u>yolov3</u>	darknet	Object Detection			
<u>yolov2</u>	darknet	Object Detection			43905
<u>BlazeFace</u>	pytorch	Classification			470

VectorBlox™ Accelerator SDK



Supports TensorFlow and ONNX (Caffe, CNTK, mxnet, PyTorch etc.)

VectorBlox™ Accelerator SDK




CNNs are supported with tutorials



Smart Camera
Reference Design




VectorBlox™ Accelerator SDK



Switch CNNs without “Re-programming”
or “Place and Route”


```
struct model_desc_t models [] = {  
    { "MobileNet V1", 0x259000, 224, IMAGENET, 10, 0x6926fe18, 30},  
    { "Tiny Yolo V3 COCO", 0xed5000, 416, YOLOV3, 30, 0xf640b794, 30},  
};  
vbx_cnn_t* the_vbx_cnn;  
int logo_width = 420;  
int logo_height = 264;  
unit32_t* logo = 0;  
  
void clear_frame (uint32_t* draw_frame) {  
    //clear relevant sections.  
    //clear center  
    draw_rectangle ((1920-1080)/2,0,1080, 1080, 0,
```



MICROCHIP


CoreVectorBlox Neural Network Engine

- **Multi-core Software Architecture**
 - Mi-V runs complete software-based application
- **VectorBlox™ Matrix Processor (MXP)**
 - Elementwise tensor operations
(add, sub, xor, shift, mul, dotprod ...)
 - Up to 8 32-bit ALUs
 - Mixing precisions ok (int8, int16, int32)
 - 256KB Scratchpad Memory and DMA Controller
- **VectorBlox CNN**
 - Tensor multiply-accumulate operations
 - Fixed at either int8 precision
 - Layer enhancements added with software update



Overlay implementation allows several different networks to run on the same FPGA design without the need to resynthesize

CoreVectorBlox IP



Configuration	Vector Processor Width	Vector Scratchpad	CNN Accelerator Array Size
V250	128-bit	64kB	16x16
V500	256-bit	128kB	16x32
V1000	256-bit	256kB	32x32


- **Control Registers**
 - Control, Status and Error signaling
- **Microcontroller**
 - Parses network structures from BLOBs
 - Controls MXP Vector Processor
- **MXP Vector Processor**
 - Processes general neural network layers and offloads convolution layers to CNN Accelerator
- **CNN Accelerator**
 - Processes convolution network layers

CoreVectorBlox v1.0 Utilization and Performance

Resources	V250	V500	V1000
Peak GOPs	79	146	279
ResNet-50 (fps)	4	8	14
Tiny Yolo v3 (fps)	5	10	28
MobileNet v1 (fps)	19	38	81
MobileNet v2 (fps)	17	34	76

Resource Utilization (LUT4)			
MPF100 LEs	25%	43%	
MPF200 LEs	14%	24%	32%
MPF300 LEs	9%	16%	21%
MPF500 LEs	6%	10%	13%

Fabric Clock	153 MHz	143 MHz	129 MHz
--------------	---------	---------	---------



Configuration Options

Size Configuration: V250, V500, V1000

Memory Master Width: 256b, 512b

- Configurable performance based on size and power budget
- V2000 (~500 GOPs), V4000 (~750GOPs) are in development

VectorBlox™ IP Performance and Utilization for Space applications

Without TMR – default implementation


Configuration	Size (no TMR)			Fmax (base) MHz	Peak GOPs	mW/GOP	Performance (FPS)	
	kLUTs	MPF500T	RTPF500T				Mobilenet-v1	TinyYOLO-v3
V250	28	6%	6%	154	79	7.0	26.2	9.1
V500	48	10%	10%	143	146	6.4	47.7	16.6
V1000	63	13%	13%	136	279	5.1	68.0	26.5

With TMR – synthesized local TMR, preliminary results

Configuration	Size (synthesized local TMR)				Fmax (base) MHz	Peak GOPs	mW/GOP	Performance (FPS)	
	kLUTs	kDFF	MPF500T	RTPF500T				Mobilenet-v1	TinyYOLO-v3
V250	56	66	14%	14%	102	52	-	17.4	6.0
V500	97	116	24%	24%	86	88	-	28.7	10.0
V1000	136	161	33%	33%	76	156	-	38.0	14.8

2-3x More Power Efficient Inferencing

Core Name	Peak GOPs	Dynamic Power (mW)	Static Power* (mW)	Total Power (mW)	Total Power (mW/GOP)
VectorBlox™ V1000	279	1094	206	1300	5.1
VectorBlox V500	146	698	127	825	6.4
VectorBlox V250	79	387	65	452	7.1
Competitor A Core 1	332.8	3072	976	4048	12.2
Competitor A Core 2	166.4	2201	528	2729	16.4





- **50% to 70% power saving for similar inferencing performance**
- **Higher inferencing performance in applications with power constraints**

*Scaled for resource utilization


5-10x Higher Performance/Watt

Source Network	Core Name	Latency (ms)	FPS	FPS/W	GFLOPS /sec
OpenVINO mobilenet-v1-1.0-224	VectorBlox™ V250	37.0	27.0	60.1	31.3
	Competitor B CNN Accelerator	530	1.9	4.8	2.2
OpenVINO mobilenet-v2	VectorBlox V250	39.0	25.6	57.0	22.9
	Competitor B CNN Accelerator	304	3.3	8.5	2.9
Competitor B MobilenetV1 detector	VectorBlox V250	7.9	126.6	281.4	18.2
	Competitor B CNN Accelerator	69	14.5	37.3	2.1
Competitor B MobilenetV2 detector	VectorBlox V250	16.3	61.3	136.3	18.8
	Competitor B CNN Accelerator	178	5.6	14.4	1.7

	Power (mW)
VectorBlox V250	450
Competitor B CNN Accelerator	388




TF Lite Support on Ver2.0 for PolarFire® Family



The VectorBlox™ Solution in PolarFire® Video

PolarFire Video Kit



```
struct model_desc_t models [] = {
    { "MobileNet V1", 0x259000, 224, IMAGENET, 10, 0x6926fe18, 30},
    { "Tiny Yolo V3 COCO", 0xed5000, 416, YOLOV3, 30, 0xf640b794, 30},
};
vbx_cnn_t* the_vbx_cnn;
int logo_width = 420;
int logo_height = 264;
uint32_t* logo = 0;

void clear_frame (uint32_t* draw_frame) {
    //clear relevant sections.
    //clear center
    draw_rectangle ((1920-1080)/2,0,1080, 1080, 0,
```

**Code in SoftConsole (Eclipse) and
write into SPI Flash**

Benchmarking



Figure 6.2: Power Consumption of FPGA in Dynamic Mode

Table 6.2: Phase[II] : MobilenetV1 build and inference run power and thermal measurements


Hardware	Microchip PolarFire FPGA Video-Kit	NVidia Jetson Nano
Input DC Power	12V	5V
OS	No OS	JetPack Linux
internal power consumption in real time	none	Jetson Nano Stats - Jtop
Inference	For 480p = 0.829A, 0.831A, For 720p = 0.834A, 0.833A, 0.834A, 0.835A, 0.837A	For 480p = 0.923A, 1.273A 1.35A, 1.695A
Heat dissipation		
Inference time	37°C maximum to 44°C	Shoots up to 57°C (at 11min:32seconds) on Thermocouple Internal GPU temperature = 63.00°C



The Jetson Nano achieves around 23FPS to 35FPS.
The PolarFire® Video Kit achieves around 37FPS to 45FPS

Does Power Matter?


Effect of Temperature on Device Reliability



- **FIT doubles with a 10° rise in junction**
- **Thermally-challenged environments (no fans, high ambient)**
 - ADAS, headlight control, PLC backplane controllers, miniature low light vision systems

What is Available Today

- **GitHub SDK**
 - Tools
 - Model Optimization
 - Calibration
 - Simulator
 - Tutorials
 - Documentation



master 1 branch 14 tags		Go to file	Code
Vectorblox Automated Publish		39af3e9 last week	20 commits
app_notes	Automated Publish	7 months ago	
docs	Automated Publish	7 months ago	
drivers/vectorblox	Automated Publish	7 months ago	
example	Automated Publish	last week	
fw	Automated Publish	last week	
lib	Automated Publish	last week	
python/vbx	Automated Publish	last week	
tutorials	Automated Publish	last week	
.gitattributes	initial commit	3 years ago	
README.md	Automated Publish	2 years ago	
install_dependencies.sh	Automated Publish	7 months ago	
install_venv.sh	Automated Publish	2 years ago	
requirements.txt	Automated Publish	7 months ago	
setup_vars.sh	Automated Publish	last year	

<https://github.com/Microchip-Vectorblox/VectorBlox-SDK>

Mi-V Ecosystem Solutions



Technolution

IP / SoftCPU

PQ Secure

PQ SHIELD

sandlogic

logictronix

ultra soc

PRAESUM COMMUNICATIONS

SiFive

Microsemi

DirectCore

Development Tools

Libero System-on-Chip

SoftConsole

Green Hills SOFTWARE

LAUTERBACH DEVELOPMENT TOOLS

SIEMENS

AdaCore Build software that matters.

MathWorks

IAR SYSTEMS

SEGGER

ASHLING

MPSI

OS / RTOS

Embedded With RTEMS

NuttX RTO

ubuntu

GNOME

seL4

NUCLEUS

ROS

μ-velOSity

amazon freeRTOS

FreeBSD

Microsoft Azure

INTEGRITY

Zephyr



Hardware

ariAS embedded Intelligence for Industry

DIGITAL CORE TECHNOLOGIES QUALITY BY DESIGN

ALDEC THE DESIGN VERIFICATION COMPANY

ENCLUSTRA FPGA SOLUTIONS

CONCLUSIVE ENGINEERING

EmCraft systems

Linera

trenz electronic

SUNDANCE DSP INC

NUMATO LAB

antmicro

CONCLUSIVE ENGINEERING

DORNERWORKS

CodeThink

trenz electronic

SUNDANCE DSP INC

ariAS embedded Intelligence for Industry

DIGITAL CORE TECHNOLOGIES QUALITY BY DESIGN

ENCLUSTRA FPGA SOLUTIONS

EMDALO TECHNOLOGIES Embedded Intelligent Systems

PRAESUM COMMUNICATIONS

OFFcode

Linera

NUMATO LAB

Middleware

antmicro

aicas

wolfSSL

Golioth

imperas

DORNERWORKS

Veridify Security

0x5 HEX-Five

YOUR EMBEDDED DESIGN SERVICES PARTNER

Design Services



Summary and Conclusion

- **VectorBlox™ Accelerator SDK creates efficient networks for FPGA deployment**
 - Optimization and conversion to 8-bit models with minimal loss of accuracy
 - Bit-accurate simulation environment allows evaluation before programming
- **CoreVectorBlox IP creates flexible, power-efficient neural networks on PolarFire and RT PolarFire® FPGAs**
 - Update and modify networks without changing the FPGA design
 - Networks can be designed and modified by engineers with no FPGA experience
- **RT PolarFire allows greater inferencing performance in space**
 - Higher GOPS/mW in power-constrained space applications
 - Radiation effects suitable for LEO, GEO and deep space
- **Start Designing Today**
 - RT PolarFire® FPGA engineering models - available today in Ceramic Land Grid, Ceramic Ball Grid and Ceramic Column Grid packages
 - Libero® development software - Libero SoC software (v12.4 and later) supports synthesized triple module redundant (TMR) flip-flops and placement constraints for SEU protection
 - PolarFire and PolarFire SoC Smart Embedded Vision Video Kits – available today

What's NEXT?




Next: The Second Generation of PolarFire® FPGAs

We're Doubling Power Efficiency... Again!


System-Level Design Gets Its Own Suite



Total 4X gain
with PolarFire 2
FPGA



Power
Efficiency



CoreMark® Score



15x

TOP/S

PolarFire SoC




PLUS...

Deterministic, asymmetric processing

HPSC-compatible core complex

Mil-grade anti-tamper and cybersecurity

Immune to single event upsets (SEU)



System Designer

- Abstracted C/C++ and RTL flows
- Complete SoC creation with C/C++
- Hardware/software partitioning and co-design
- Full SoC debug to speed bring-up
- Pre-silicon software development

FPGA Designer

- Modern and intuitive UI
- Fast timing closure
- Adaptive design flows
- Better IP handling and delivery
- Faster simulation and expanded debug

HPSC Architecture Highlights

Advanced

FinFET Process




Machine Learning



Multi-Core

RISC-V

Space Compute



Time Sensitive

Ethernet

Networking



PCIe

SWaP Optimized Heterogenous


Fault Tolerant

Architecture



Real-time Processing

Deterministic Latency



Vector Engines



Secure Enclave

Co-processor Interfaces



Extensible Power & Performance



HPSC is Critical for NASA's Strategic Framework

HPSC-based Autonomy is the Key

Go

Rapid and efficient space transportation

Rapid, low cost robotic payloads: Moon, Mars and beyond.
Next generation computing and communications in space.

Land

Expanded access to diverse surface destinations

Precision landings. Avoid unknown local hazards.
Ability to handle unknown environments and situations.

Live

Sustainable Living and Working Farther from Earth

Human and Robotic Lunar Missions > 28 days.
Human Mars Missions > 800 days.

Explore

Transformative missions and discoveries

Robotic systems augmenting human operation.
Remote servicing, assembly and manufacturing.

Questions


Appendix

Supported Layers in CoreVectorBlox IP

Supported CNN Layers		
Abs	Add	AveragePool
BatchNormalization	Concat	Conv
Dropout	Gemm	Identity
LeakyRelu	MatMul	MaxPool
Mul	PRelu	Relu

- Additional layers can be added via software update

The PolarFire® SoC Smart Embedded Vision Kit



Target Applications

- Edge Compute
- IoT
- AI/ ML
 - Post Processing
- Embedded Vision
 - OpenVX
 - Streaming over UDP
 - 1588 Time Stamping

New