



赛灵思工业物联网研讨会
XILINX IIoT SEMINAR

Machine Vision Trends & Solutions

翁羽翔 Trevor

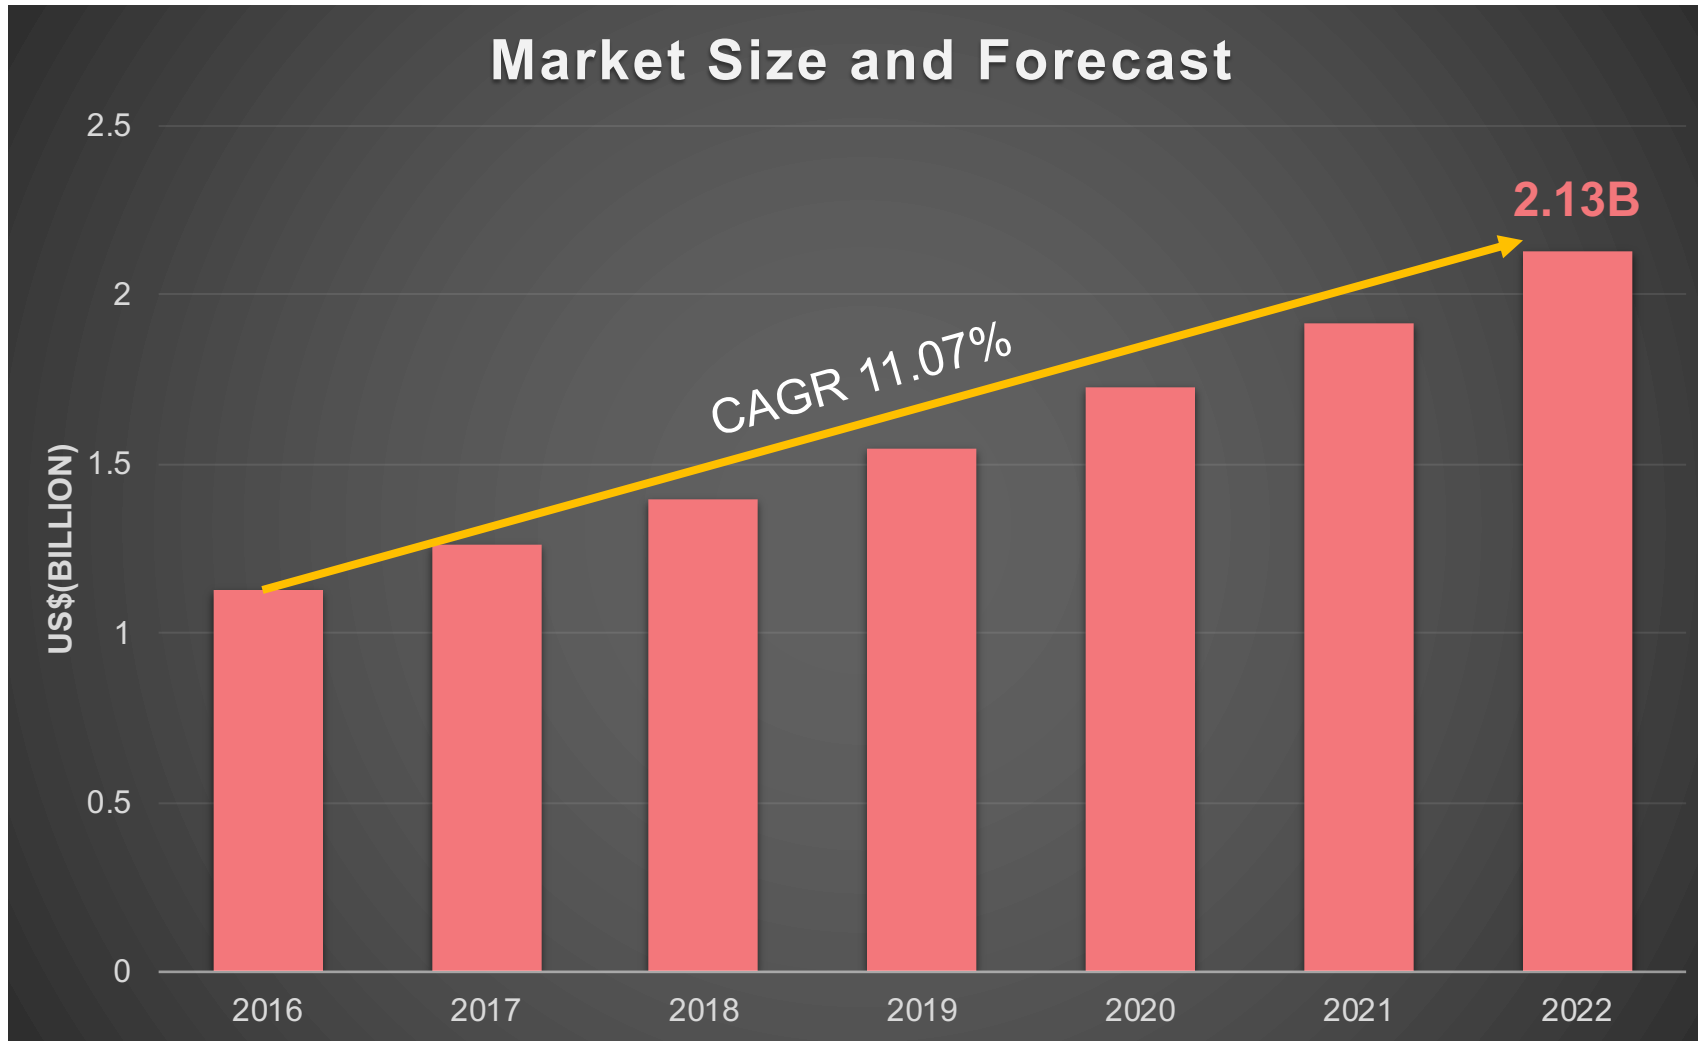
赛灵思工业与医疗市场经理

2019年5月

3D Vision in Robotics



3D Machine Vision Market Size and Forecast



Source : Markets and Markets Analysis

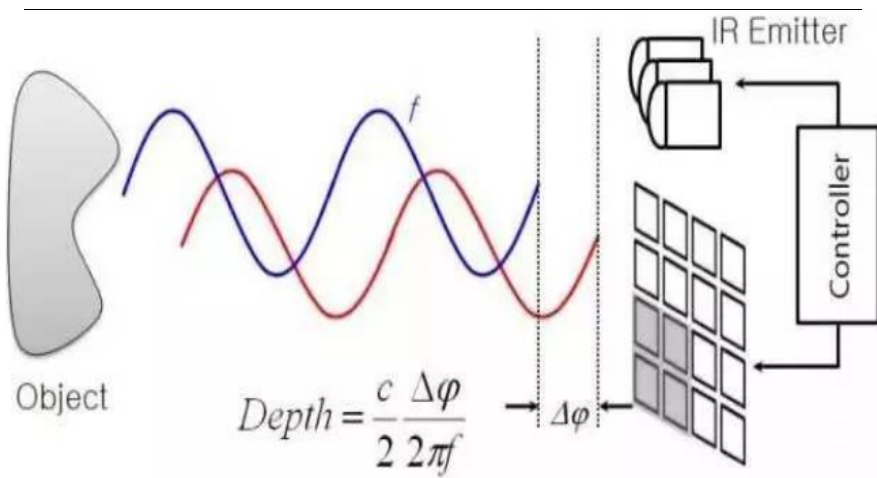
Three Key Factors of Growth:

Robotics - 3D Vision Integration


Manufacturing - 3D Object Analysis

Non-Industrial - 3D Machine Vision


3D Vision Technology



Time of Flight



Stereo Vision



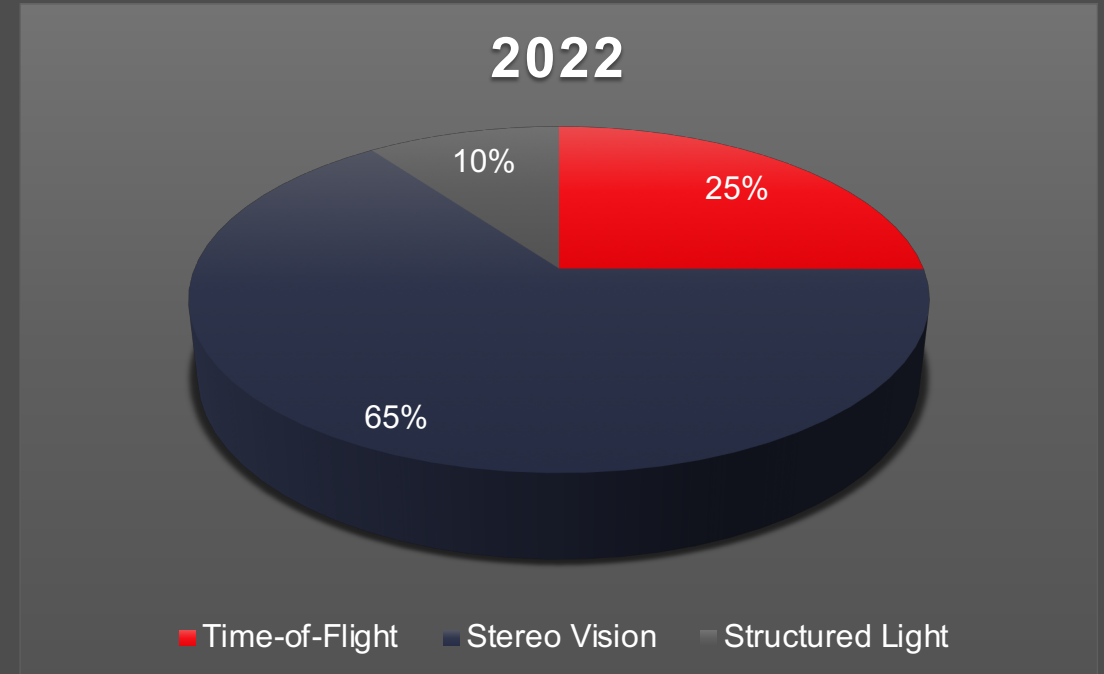
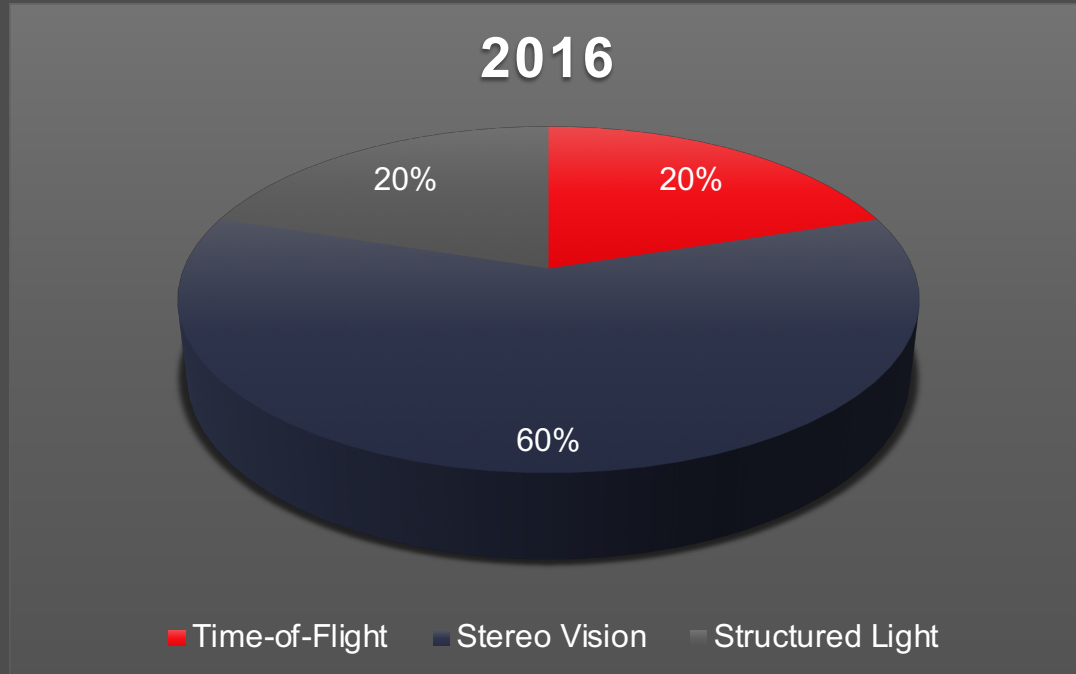
$$[x, y, z] = \frac{b}{f \cdot \text{ctg} \theta - x'} [x', y', f]$$

Structured-Light

3D Solution Comparison

类型	TOF	RGB双目	结构光
工作原理	依据光的飞行时间	RGB图像特性点匹配 三角测量间接计算	主动投射已知编码图案
测量精度	厘米级	毫米级	0.01mm-1mm
测量范围	100m以内	2m以内, 距离远, 测距不准确	10m以内
影响因素	不受光照变化和表面纹理影响	受光照变化和表面纹理影响, 夜晚无法使用	不受光照变化和表面纹理影响; 容易被强光淹没, 不适合室外
其他	分辨率一般 帧率较高 功耗较高, 全面照射	分辨率高 帧率高 功耗较低	分辨率较高 帧率一般 功耗一般, 局部照射

3D Camera Market Share and Forecast



Source : 2016 Allied Market Search Report

Seeing In the Dark

> Dual Sensor (IR + Visible Light)




> Enhanced ISP

Brightness Full Color In the Dark Colourfulness



Burst of raw frames

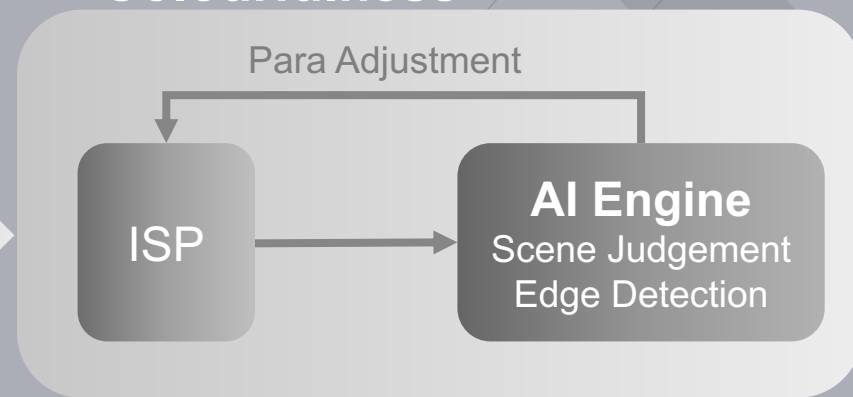


Merged raw image



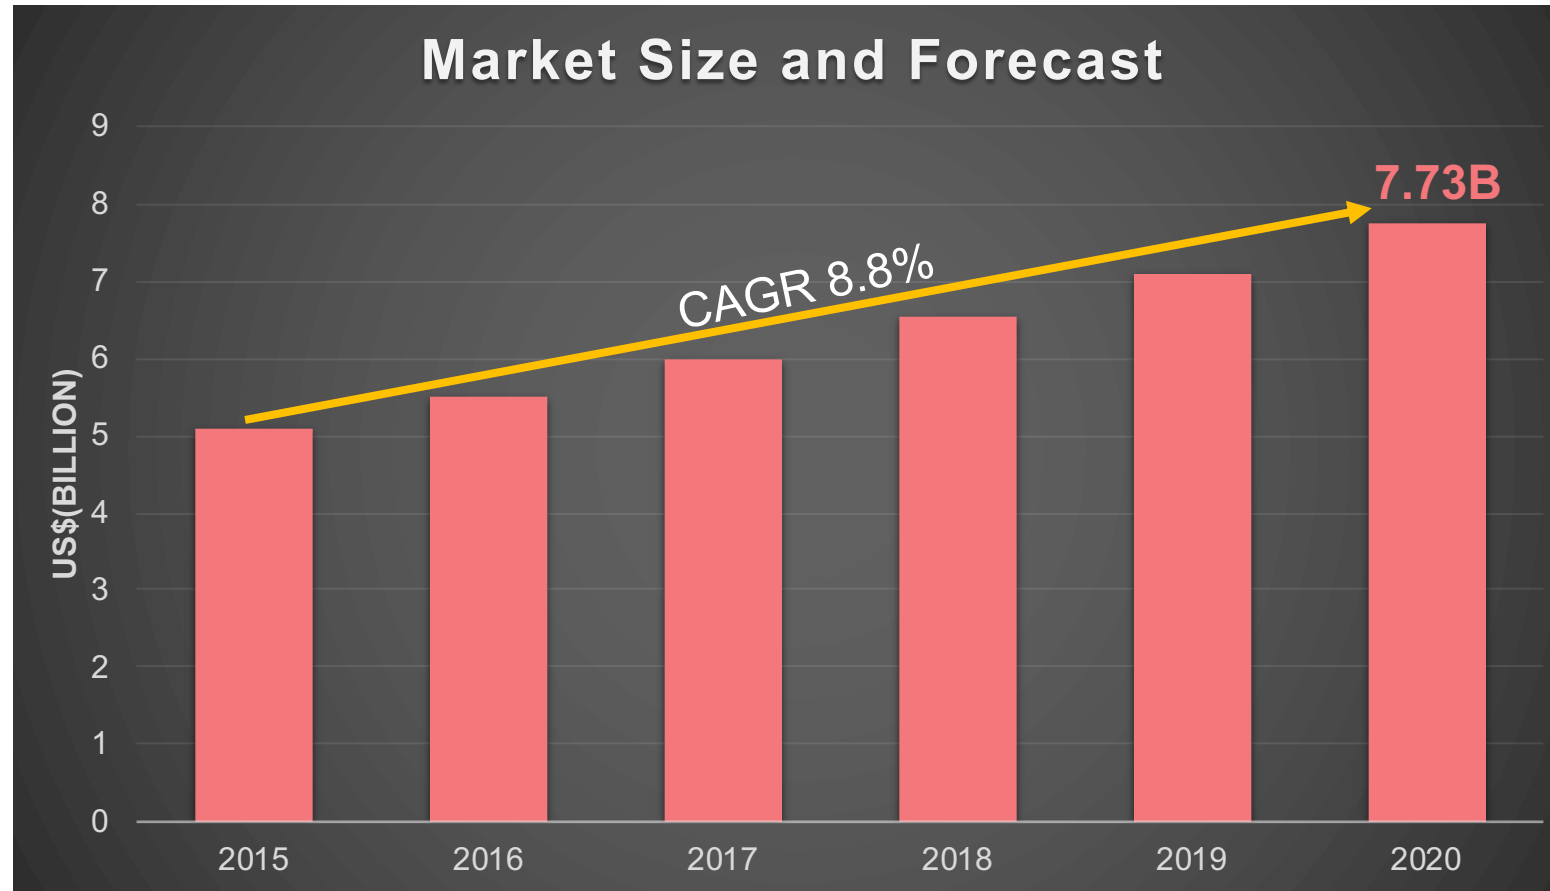
Final high-quality result

HDR+



AI + ISP

Night Vision Market Overview and Forecast



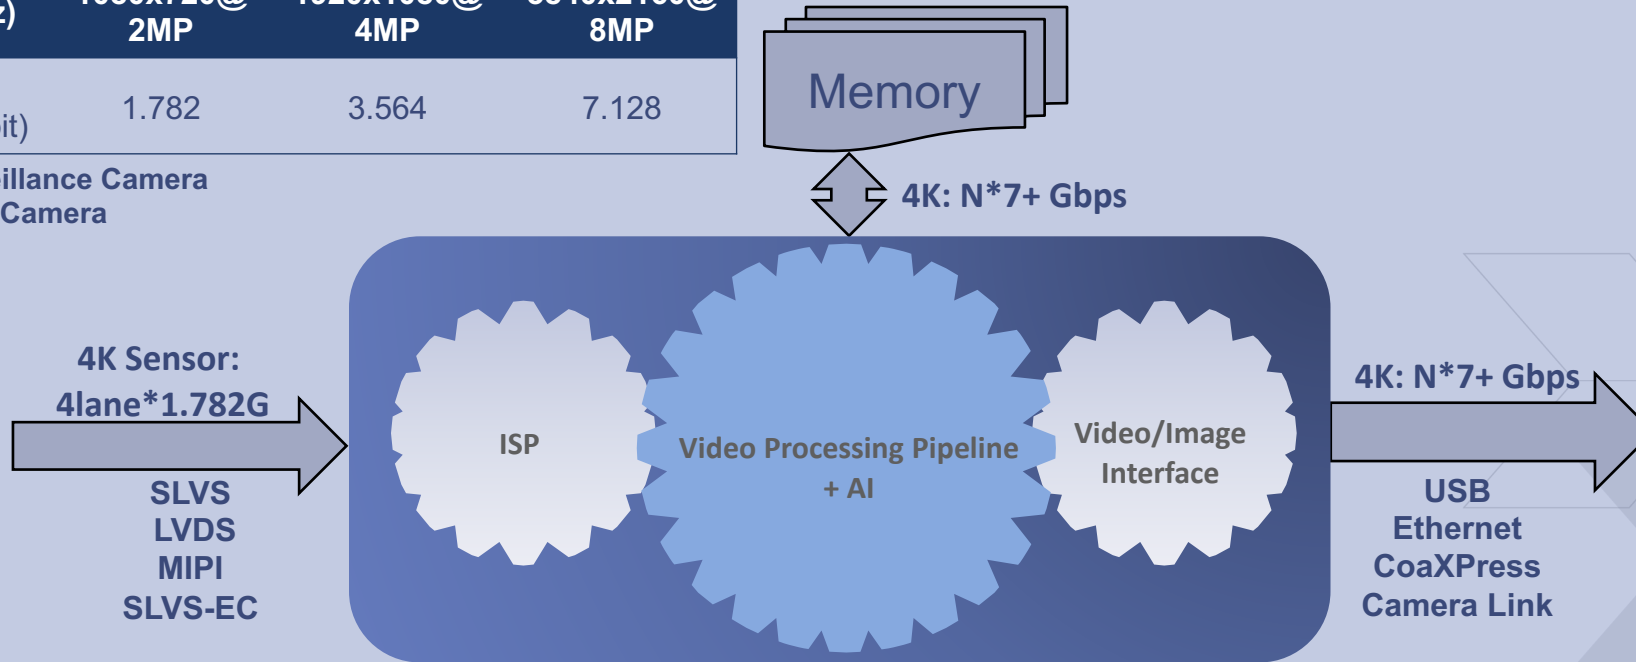
Source : Markets and Markets Analysis

Industrial Vision Evolutions

> High-Resolution Sensor, High-Speed I/F, High-Throughput System

Resolution(60Hz)	1080x720@ 2MP	1920x1080@ 4MP	3840x2160@ 8MP
Bit Rate(Gbps)			
Bits Per pixel (12bit)	1.782	3.564	7.128


8MP for Video Surveillance Camera
8MP to 29MP for FA Camera



No ASSPs / DSPs to Effectively Address this Challenge

© Copyright 2019 Xilinx

Onsemi – Xilinx Joint Demo in 2019 China Vision Show

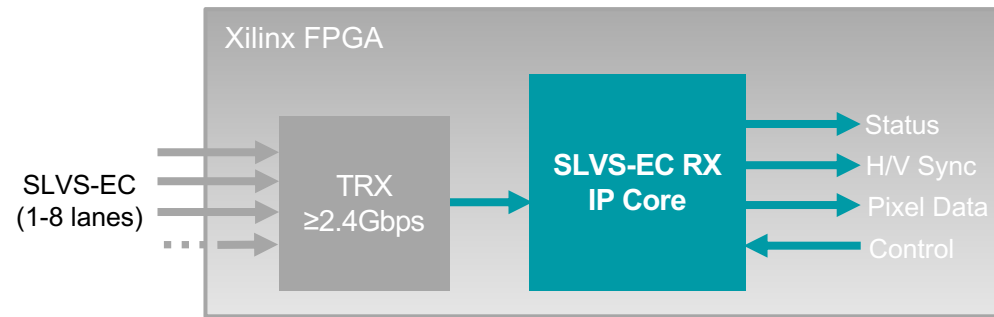


ZCU104 Demo for HiSPi
XGS 12000 – 4K CMOS Global Shutter

Framos - Sony - Xilinx

> SLVS-EC RX IP Core for Xilinx FPGAs

- >> Receiver IP Core
- >> Supporting SLVS-EC v1.2
- >> AXI4-Lite control interface
- >> 1, 2, 4, 8 Lane Support
 - Configurable by user
- >> Full interface speed (2.376Gbps)
 - Baud Grade 2
- >> All supported pixel formats
 - RAW: 8, 10, 12, 14 bits per pixel
 - Dynamic mode change
- >> Error detection
 - CRC – Cyclic Redundancy Check
- >> Embedded data extraction

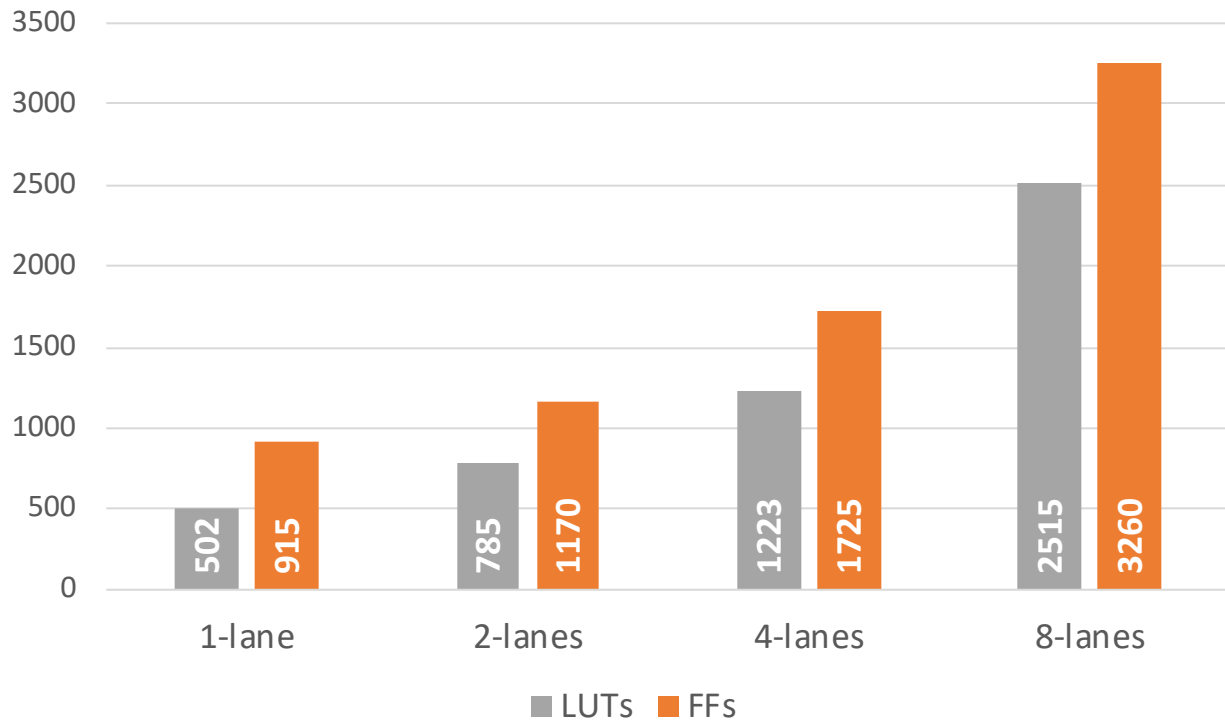


> Supported Device Family


- >> Xilinx Artix-7™
- >> Xilinx Kintex-7™
- >> Xilinx Zynq-7000™ SoC
- >> Xilinx Kintex Ultrascale™
- >> Xilinx Kintex Ultrascale+™
- >> Xilinx Zynq UltraScale+™ MPSoC

Could be More, Just Raise Your Hand

Framos SLVS-EC IP Resource & Demonstration



SLVS-EC IP Resource



SLVS-EC IP Eval Kit

High-Speed I/F (3D Vision, Higher Resolution & Faster Frames)

GigE 1.2 -> GigE 2.0 -> GigE 2.1

CXP 1.1 -> CXP 2.0 -> CXP 2.1

USB3.0 -> USB3.1 -> USB3.2

Camera Link HS -> CLHS XP
(CLHS)

SLVS-EC 1.2, MIPI 2.1

GigE -> 2.5/5GigE(NBASE-T) -> 10GigE

CXP-6 -> CXP-10/12

5Gbps -> 10Gbps -> 20Gbps

2.4Gbps -> 10Gbps

8K @ 30fps

Not Only Data Rate, But Also IEEE 1588, GenDC(for 3D Data Format) ...

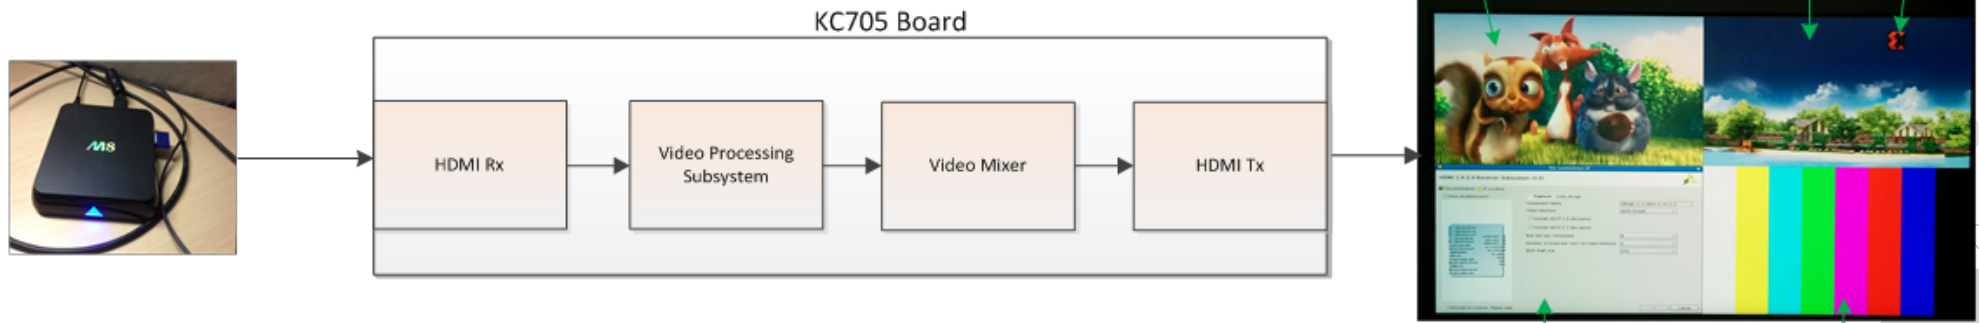
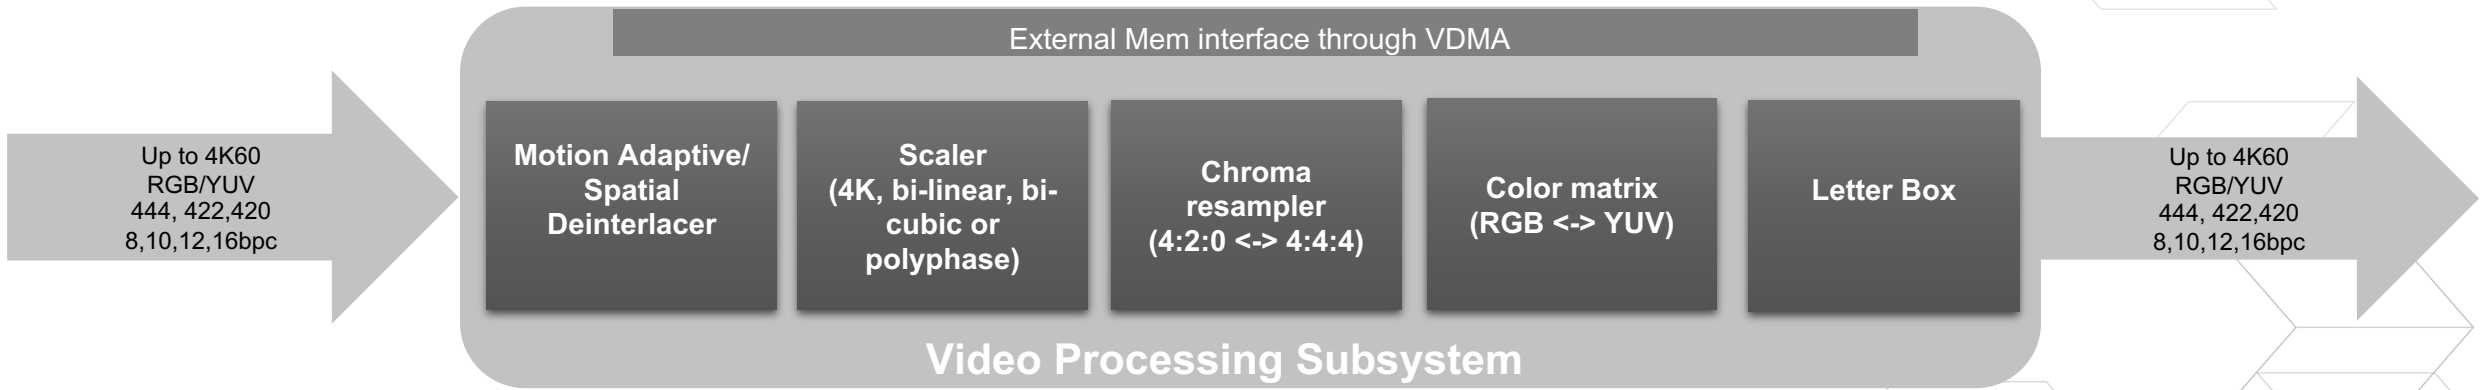
Xilinx MIPI Solution – Higher Data Rates Rollout

	28nm Family	Kintex / Virtex UltraScale	UltraScale+
MIPI D-PHY	External D-PHY chip 800Mbps/Lane	External D-PHY chip 800Mbps/Lane	Native IO support 1.5Gbps/Lane
MIPI DSI / Host (TX)	Xilinx	NWL	Xilinx
MIPI CSI2 / RX	Xilinx	NWL	Xilinx
MIPI CSI2 / TX	Xilinx	NWL	Xilinx

US+*
2.5Gbps/Lane
In 2019.1

*2.5Gbps support on all US+ devices with exception of A484 package devices

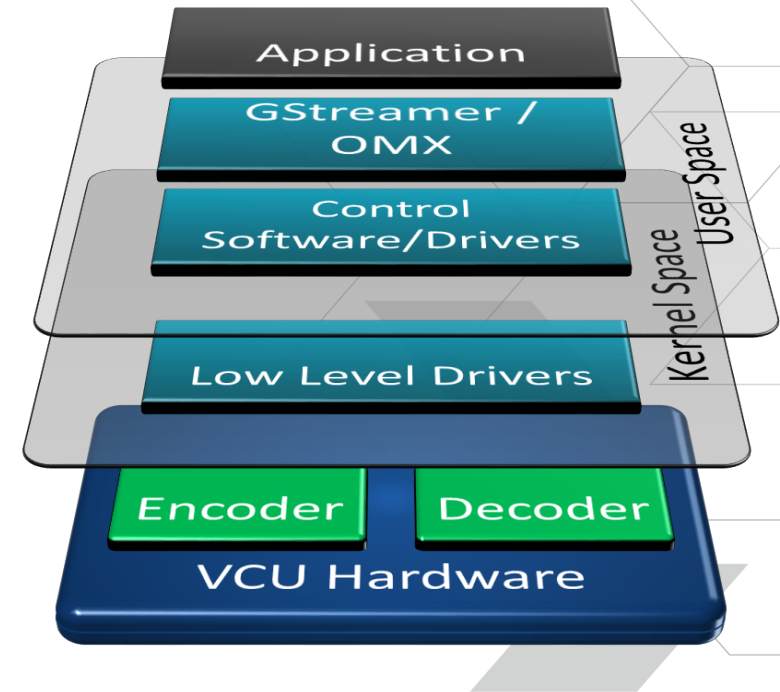
Xilinx Video Processing Subsystem (VPSS)



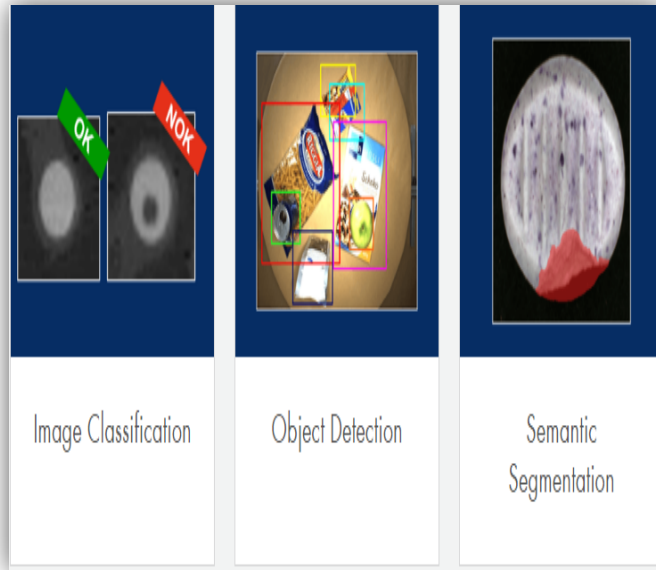
XAPP1291 – VPSS Ref Design
UP/Down/Cross conversion up to 4K using VPSS
Video Mixer with 4 layer mixing and logo insertion
Run-time configuration through SW application

Coming Out in 2019


- > HDMI2.1 -> 8K @ 60fps
- > DP1.4 -> 8K @ 60fps
- > VPSS -> 8K @ 30fps
- > **Hardened Video Codec Unit (OMX & Control Software)**
 - >> Low Latency (Encoder + Decoder)
~1 Frame(1/60fps)
 - >> PL DDR Extension to 4K @ 60fps/8K @ 15fps
 - >> 32 Streams 480P



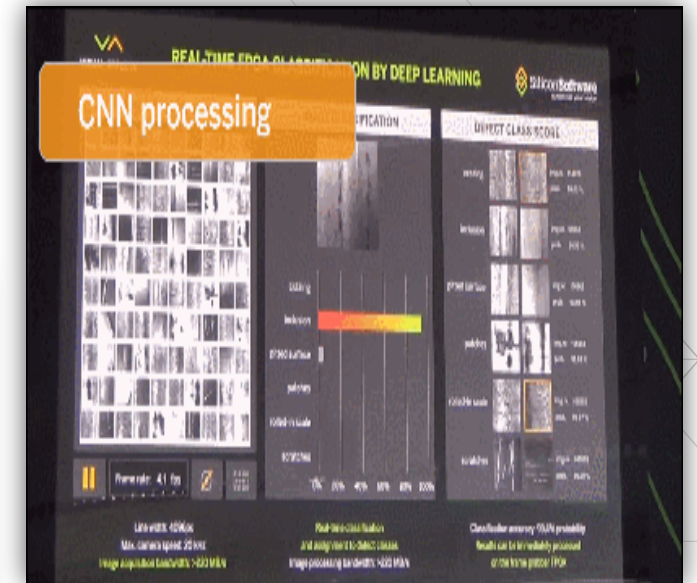
Deep Learning in Machine Vision (Xilinx Partners)



MVTec - HALCON



IDS_Kortiq



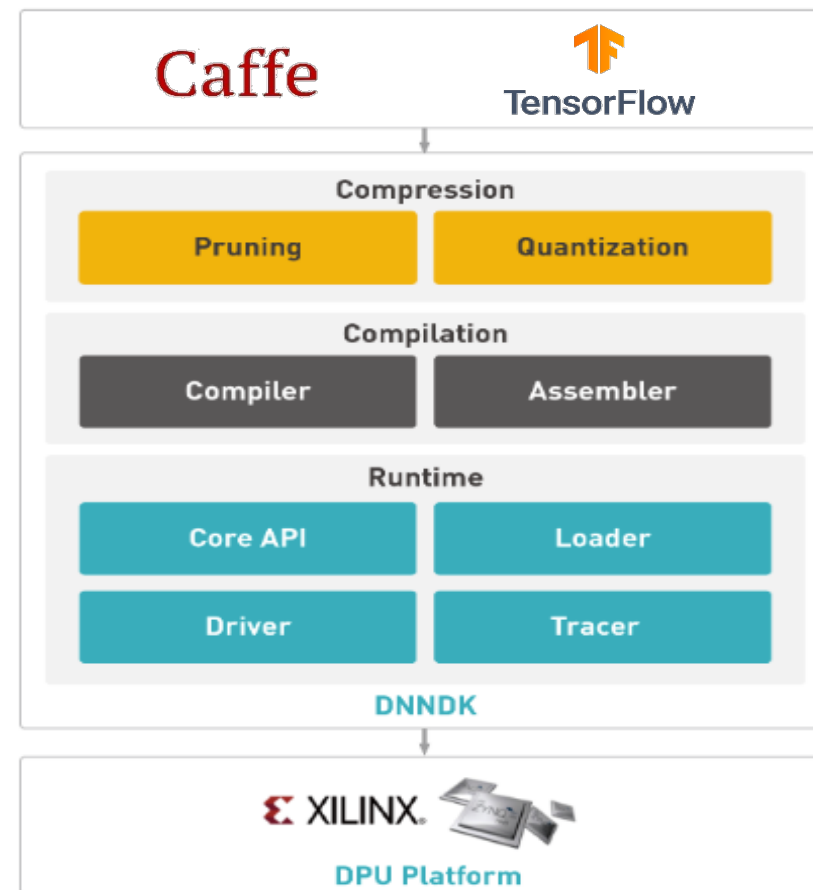
Silicon Software - VisualApplets

Xilinx ML Stack



reVISION Stack

Xilinx DNNDK v3.0 Enables TensorFlow



Xilinx - Deephi Stack




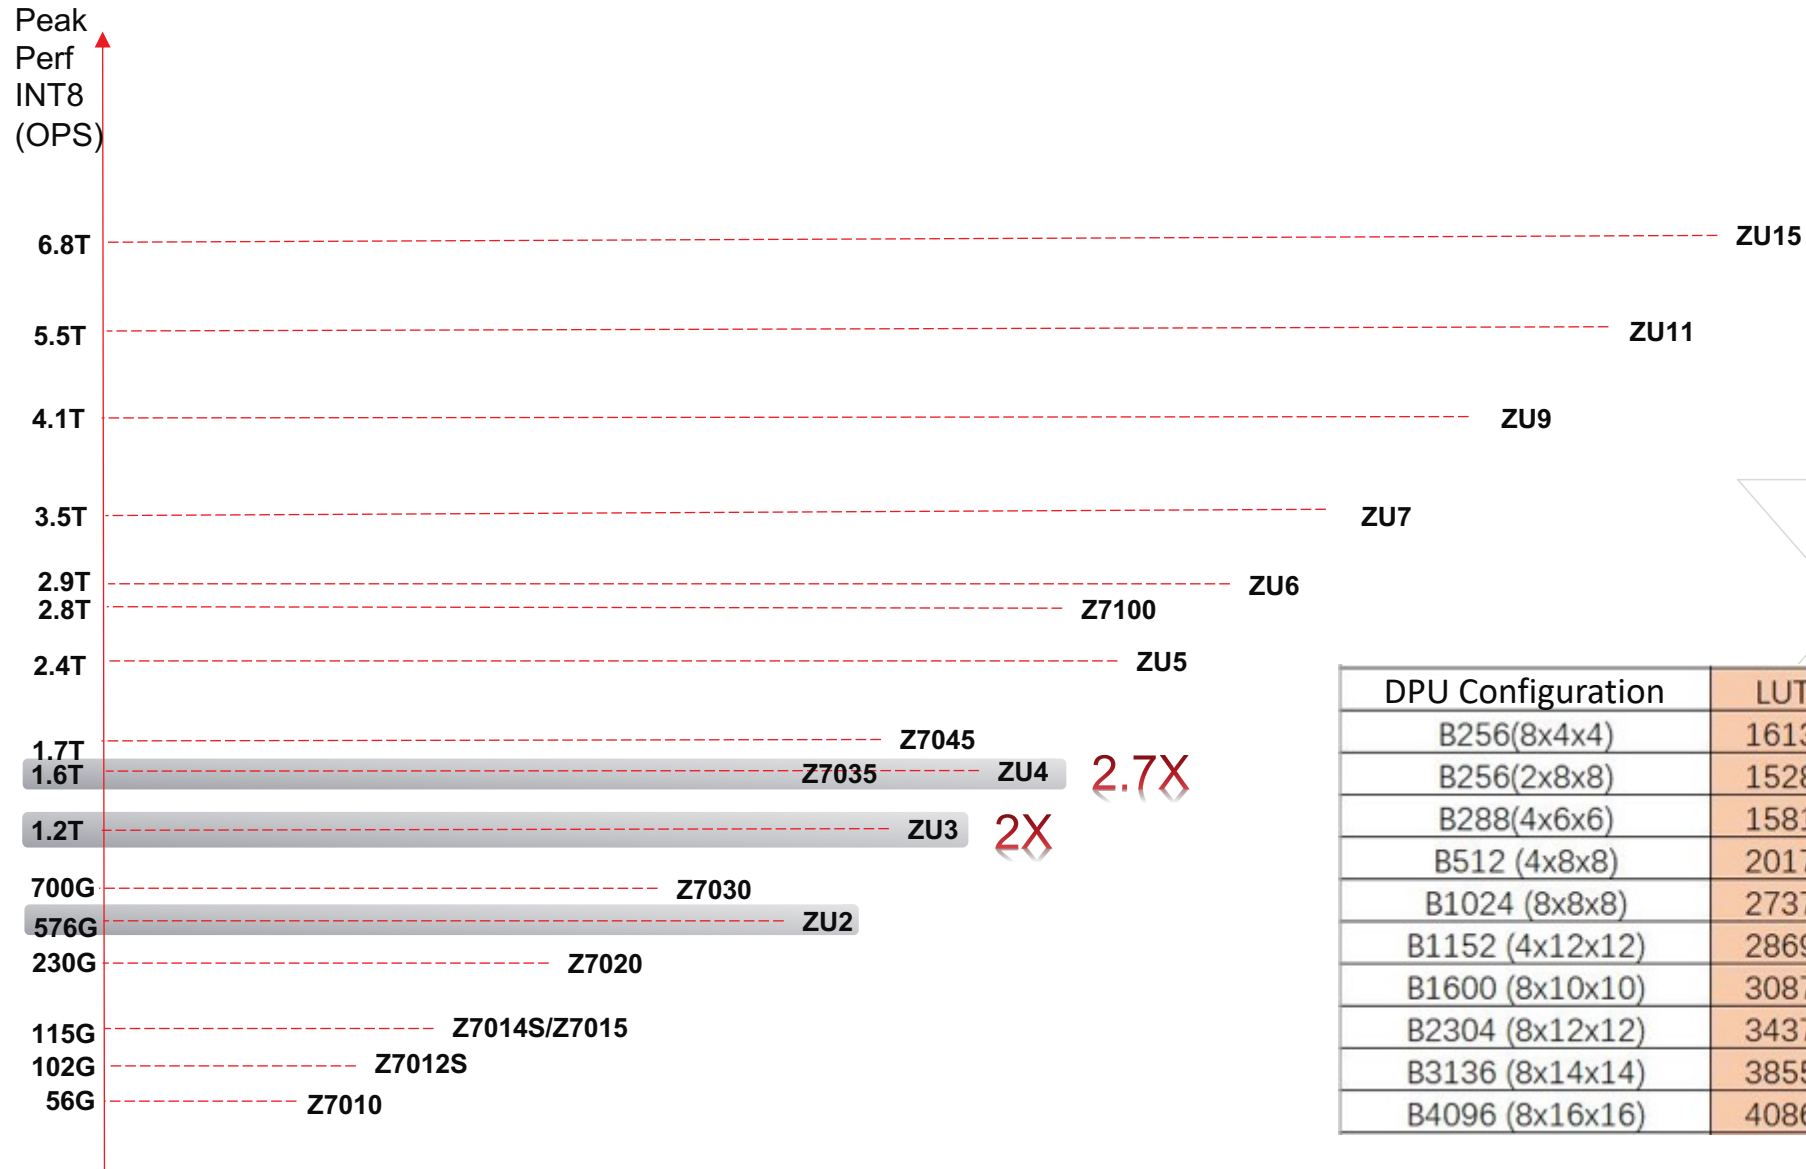
xfOpenCV Library

Basic Functionality	Geometric Transforms	Image Processing and Filters	Feature Detection and Classifiers	3D Reconstruction	Motion Analysis and Tracking
Absolute difference	Scale/Resize	Box	Canny edge detection	StereoLBM	Mean Shift Tracking (MST)
Accumulate	StereoRectify	Gaussian	Fast corner	StereoSGM	LK Dense Optical Flow
Accumulate squared	Warp Affine	Median	SVM (binary)		
Accumulate weighted	Warp Perspective	Sobel	Harris corner		
Arithmetic addition	Remap	Custom convolution	Histogram of Oriented Gradients (HOG)		
Arithmetic subtraction		Equalize Histogram	Houghlines		
Bitwise: AND, OR, XOR, NOT		Dilate			
Pixel-wise multiplication		Erode			
Channel combine		Bilateral			
Channel extract		OTSU Thresholding			
Color convert		Thresholding			
Convert bit depth		Image pyramid			
Table lookup		Color Detection			
		Integral image			
		Gradient Magnitude			
		Histogram			
		Gradient Phase			
		Min/Max Location			
		Mean & Standard Deviation			

xfOpenCV Library

Application	Module	Algorithm	Model Development	Compression	Deployment
Face	Face detection	SSD, Densebox	✓	✓	✓
	Landmark Localization	Coordinates Regression	✓	N / A	✓
	Face recognition	ResNet + Triplet / A-softmax Loss	✓	✓	✓
	Face attributes recognition	Classification and regression	✓	N / A	✓
Pedestrian	Pedestrian Detection	SSD	✓	✓	✓
	Pose Estimation	Coordinates Regression	✓	✓	✓
	Person Re-identification	ResNet + Loss Fusion	✓		
Video Analytics	Object detection	SSD, RefineDet	✓	✓	✓
	Pedestrian Attributes Recognition	GoogleNet	✓	✓	✓
	Car Attributes Recognition	GoogleNet	✓	✓	✓
	Car Logo Detection	DenseBox	✓	✓	
	Car Logo Recognition	GoogleNet + Loss Fusion	✓	✓	
	License Plate Detection	Modified DenseBox	✓	✓	✓
	License Plate Recognition	GoogleNet + Multi-task Learning	✓	✓	✓
ADAS/AD	Object Detection	SSD, YOLOv2, YOLOv3	✓	✓	✓
	3D Car Detection	F-PointNet, AVOD-FPN	✓		
	Lane Detection	VPGNet	✓	✓	✓
	Traffic Sign Detection	Modified SSD	✓		
	Semantic Segmentation	FPN	✓	✓	✓
	Drivable Space Detection	MobilenetV2-FPN	✓		
	Multi-task (Detection+Segmentation)	DeepHi	✓		


DPU Scalability




DPU Configuration	LUTs	Registers	BRAM	DSP
B256(8x4x4)	16132	25064	43	66
B256(2x8x8)	15286	22624	53.5	50
B288(4x6x6)	15812	23689	46	62
B512 (4x8x8)	20177	31782	69.5	98
B1024 (8x8x8)	27377	46241	101.5	194
B1152 (4x12x12)	28698	46906	117.5	194
B1600 (8x10x10)	30877	56267	123	282
B2304 (8x12x12)	34379	67481	161.5	386
B3136 (8x14x14)	38555	79867	203.5	506
B4096 (8x16x16)	40865	92630	249.5	642

* B256/288/512/3136 work in progress

What about NANO?



128 Core Maxwell
472GFLOPs(FP16)
Video Encode (4K@30)
Video Decode (4K@60)
JETSON NANO



XCZU2EG-2SFVA625I
576GFLOPs(INT8)
DP-8020

DL Computing Capability

GoogleNet (inception v1)

```
xx@jetson:/usr/src/tensorrt/bin$ vim ../data/googlenet/googlenet.prototxt
xx@jetson:/usr/src/tensorrt/bin$ ./trtexec --output=prob --deploy=../data/googlenet/googlenet.prototxt --
fp16 --batch=1
output: prob
deploy: ../data/googlenet/googlenet.prototxt
fp16
batch: 1
Input "data": 3x224x224
Output "prob": 1000x1x1
name=data, bindingIndex=0, buffers.size()=2
name=prob, bindingIndex=1, buffers.size()=2
Average over 10 runs is 12.0836 ms (host walltime is 12.14)
Average over 10 runs is 12.0969 ms (host walltime is 12.14)
Average over 10 runs is 12.0677 ms (host walltime is 12.11)
```

82 fps

ResNet-50

```
deploy: ../data/googlenet/ResNet50_224x224.prototxt
Average over 10 runs is 26.9436 ms (host walltime is 27.00)
Average over 10 runs is 27.3821 ms (host walltime is 27.43)
Average over 10 runs is 26.9322 ms (host walltime is 26.98)
Average over 10 runs is 27.4419 ms (host walltime is 27.49)
```

37 fps

Int8 support requested on hardware without native Int8 support, performance will be negatively affected

```
Average over 10 runs is 48.4162 ms (host walltime is 48.43)
Average over 10 runs is 49.337 ms (host walltime is 49.40)
Average over 10 runs is 48.377 ms (host walltime is 48.43)
```

20 fps

```
root@dp-8020:~# dexplorer -w
[DPU IP Spec]
IP Timestamp   : 2018-11-15 10:00:00
DPU Core Count : 1
IRQ Base 0    : 121
IRQ Base 1    : 136

[DPU Core List]
DPU Core      : #0
DPU Enabled   : Yes
DPU Arch      : B1152F
DPU Target    : v1.3.7
DPU Frequency : 430 MHz
DPU IRQ       : 138
DPU Features  : Avg-Pooling, Leaky ReLU, Depthwise Conv
```


DPU Configuration in DP8020

Neural Network	MAC (GOPS)	fps
Inception-v1	3.2	79
ResNet-50	7.7	35
MobileNet	0.56	162

UG1327 – DNNDK User Guide

DP8020 Performance w/o Pruning

NANO ≈ DP8020



Adaptable.
Intelligent.



赛灵思工业物联网研讨会
XILINX IIoT SEMINAR

 XILINX®