FTF FREESCALE TECHNOLOGY FORUM 2014

Leverage the **i.MX 6 Series Applications Processor** for Industrial Applications

FTF-IND-F0244

Hans-Peter Rosinger | FAE

APR.2014

External Use

Session Introduction

- This session talks about how to maximize computing performance and take advantage of flexible interfaces using the i.MX 6 series applications processor
  - Learn how to enable and use common acceleration units on i.MX 6 series
  - Recognize the benefit and flexibility i.MX 6 series offers for attaching peripherals
- About myself:
  - My name is Hans-Peter Rosinger, I am an Freescale field application engineer for the i.MX processor and Vybrid controller solution product lines.
  - I am based in Munich, Germany, and support industrial and general embedded customers.
- This session will last about 50min with Q&A at the end.

Session Objectives

- After completing this session you will be able to:
  - Specify industrial grade i.MX 6 series processors which are pin compatible and part of the Freescale product longevity program
  - Describe acceleration units the i.MX 6 series processor offers
  - Enable and use them on i.MX 6 series processors running Linux
  - Select the most appropriate interface on the i.MX 6 series for connecting peripherals depending on your application requirements

Agenda

- Overview of i.MX 6 series processors
  - Target applications
  - Features and industrial highlights
- Maximize i.MX 6 series computing performance
  - Take advantage of the NEON on chip acceleration unit
  - Integrate it into your application software
  - Let the GPU do the job featuring OpenCL
  - Power vs. performance trade off
- Connect peripherals to i.MX 6 series
  - Leverage many feature rich and flexible interfaces
  - Low latency / high throughput options on i.MX 6 series connecting an ASIC or FPGA
- Summary

i.MX 6 Series Target Applications

Automotive

- Infotainment
- Telematics
- Instrument Clusters
- Vision/Camera Systems

eReaders

- Monochrome eReader
- Color eReaders

Smart Devices

- IPTV/Streaming Media
- Transportation – in-flight infotainment, marine navigation, in-car signage
- Health Care – patient monitoring, telehealth, fitness equipment
- Factory, process and building automation
- Handheld scanners and printers
- Digital signage

- Thin Client and cloud server
- Point of sale terminal and printer
- Tablets – medical, educational, industrial
- IPTV/Streaming media
- IP phone
- Test and Measurement
- Appliances
- Home audio

i.MX 6 Series Industrial Grade Features

Features vary by product family

- **ARM® Cortex™-A9 based solutions** with up to 800MHz
- **Video playback in high definition** for stunning visual experience
- **Triple Play graphics architecture** enables graphics with separate 3D and 2D graphics engines
- Display interfaces **LVDS, Parallel and HDMI** display technology
- **Optimized peripheral sets:** PCIe/SATA/GbE, USB/SD/MIPI and CAN/MLB150 tailored for the auto infotainment, consumer and industrial markets
- **64-bit memory bus for optimal performance,** and memory interfaces support DDR3, DDR3L and LP-DDR2
- **Integrated power management** simplifies design and reduces external components
- **Specialized support for harsh conditions** unique to automotive and industrial environments (i.e. extreme temperatures)

i.MX 6 Series: One Platform, Differentiated Products

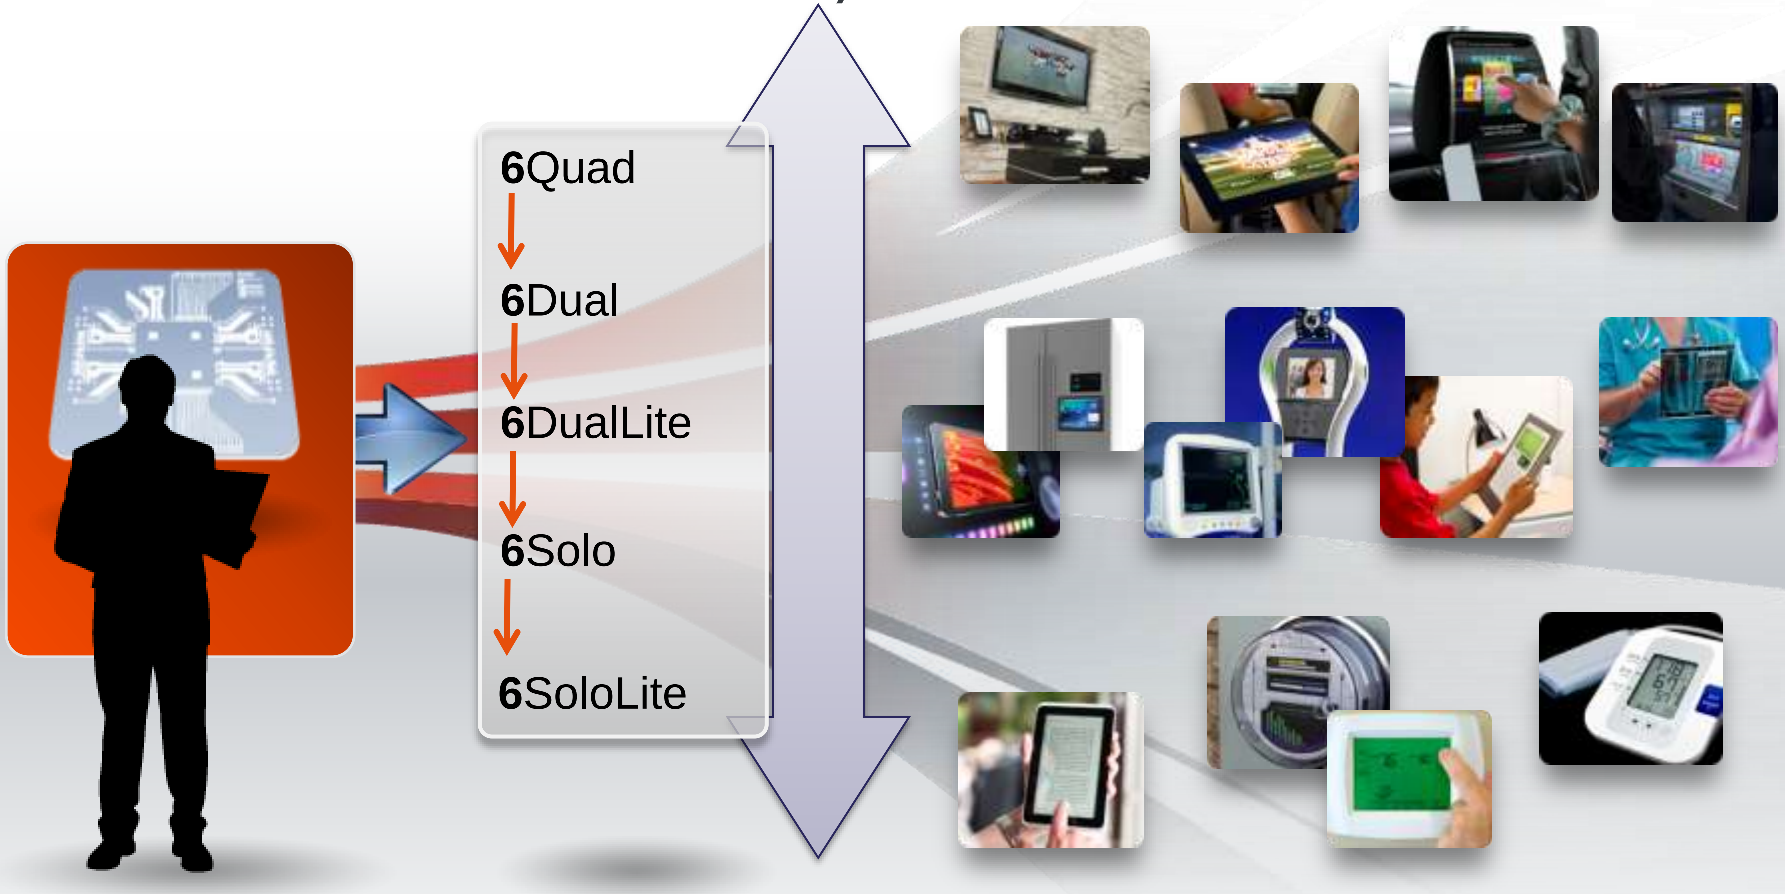

6Quad

6Dual

6DualLite

6Solo

6SoloLite

- Provide customers and partners with the **broadest range** of choices
- Reduce **development costs** and **improve** time to market
- **Performance scalability** is key to implement this strategy

Red indicates change from column to the left

6 Series At a Glance

Scalable series of five ARM Cortex A9-based SoC families

i.MX 6SoloLite

- Single ARM® Cortex™-A9 at 1.0GHz
- 256KB L2 cache, Neon, VFPvd16, Trustzone
- 2D graphics
- 32-bit DDR3 and LPDDR2 at 400MHz
- Integrated EPD controller

i.MX 6Solo

- Single ARM Cortex-A9 at 1.0GHz
- **512KB** L2 cache, Neon, VFPvd16, Trustzone
- **3D graphics** with 1 shader
- 2D graphics
- 32-bit DDR3 and LPDDR2 at 400MHz
- Integrated EPD controller

i.MX 6DualLite

- **Dual** ARM Cortex-A9 at 1.0GHz
- 512KB L2 cache, Neon, VFPvd16, Trustzone
- 3D graphics with 1 shader
- 2D graphics
- **64-bit** DDR3 and 2-channel 32-bit LPDDR2 at 400MHz
- Integrated EPD controller

i.MX 6Dual

- **Dual** ARM Cortex-A9 at 1/**1.2GHz**
- **1 MB** L2 cache, Neon, VFPvd16, Trustzone
- 3D graphics with **4 shaders**
- **Two** 2D graphics engines
- 64-bit DDR3 and 2-channel 32-bit LPDDR2 at **533MHz**
- Integrated **SATA-II**

i.MX 6Quad

- **Quad** ARM Cortex-A9 at 1.2GHz
- 1 MB L2 cache, Neon, VFPvd16, Trustzone
- 3D graphics with 4 shaders
- Two 2D graphics engines
- 64-bit DDR3 and 2-channel 32-bit LPDDR2 at 533MHz
- Integrated SATA-II

Pin-to-pin and Power Compatible

Software Compatible

i.MX 6 Series Highlights

- ARM Cortex-A9 based solutions ranging up to 1.2GHz
- HD 1080p encode and decode (except 6SL)
- 3D video playback in High definition (except 6SL)
- SW support: Google Android™, Windows® Embedded CE, Linux®

Features vary by product family

i.MX 6 Series for Industrial

- Broad Software Support
  - Mature Freescale Linux offering with LTS Kernel support
  - Broad Android offering with 3.x and 4.x
  - Windows Embedded support through proven partners
  - RTOS support including QNX, VxWorks, Integrity
  - Strong Ecosystem with Module Manufacturers and Software Integrators
- Industrial Environment and Peripheral Support
  - 3.3V IO
  - 0.8mm pitch package to reduce PCB & manufacturing cost
  - Flexible interface supporting LP-DDR2, DDR3 and LV-DDR3
  - Broad peripheral support – Gbit ethernet, CAN, SATA, PCI-E, LVDS, MIPI
  - Extended Temp range available with full Industrial qualified parts
- Up to 15 year lifetime support for manufacturing

i.MX 6 Series Product Longevity

- Certain devices in the i.MX 6 series have been added to Freescale's product longevity program
- Visit freescale.com/productlongevity for more information

i.MX 6 Family	Part #	Description	Package	Temp Range	Longevity	Product Launch
i.MX 6Quad	MCIMX6Q6AVT10AD	1GHz Automotive	21x21, 0.8mm FCBGA lidded	-40C to +125C	15yrs	2012-11
i.MX 6Quad	MCIMX6Q7CVT08AD	800MHz Industrial	21x21, 0.8mm FCBGA lidded	-40C to +105C	10yrs	2012-11
i.MX 6Dual	MCIMX6D6AVT10AD	1GHz Automotive	21x21, 0.8mm FCBGA lidded	-40C to +125C	15yrs	2012-11
i.MX 6Dual	MCIMX6D7CVT08AD	800MHz Industrial	21x21, 0.8mm FCBGA lidded	-40C to +105C	10yrs	2012-11
i.MX 6DualLite	MCIMX6U6AVM08AC	800MHz Automotive	21x21, 0.8mm MAPBGA	-40C to +125C	15yrs	2012-11
i.MX 6DualLite	MCIMX6U7CVM08AC	800MHz Industrial	21x21, 0.8mm MAPBGA	-40C to +105C	10yrs	2012-11
i.MX 6Solo	MCIMX6S6AVM08AC	800MHz Automotive	21x21, 0.8mm MAPBGA	-40C to +125C	15yrs	2012-11
i.MX 6Solo	MCIMX6S7CVM08AC	800MHz Industrial	21x21, 0.8mm MAPBGA	-40C to +105C	10yrs	2012-11

Maximize i.MX 6 Series Computing Performance

Why Computing Performance in Industrial?

- Application for object detection / gesture recognition
  - Require mathematical algorithms
- Growing display and camera resolution required for HMI
  - Image processing capabilities
- Demand for similar “consumer like” user experience
  - Graphics acceleration
- Battery powered devices in industrial automation
  - Compute smart and power efficient

i.MX 6Quad/6Dual Applications Processor

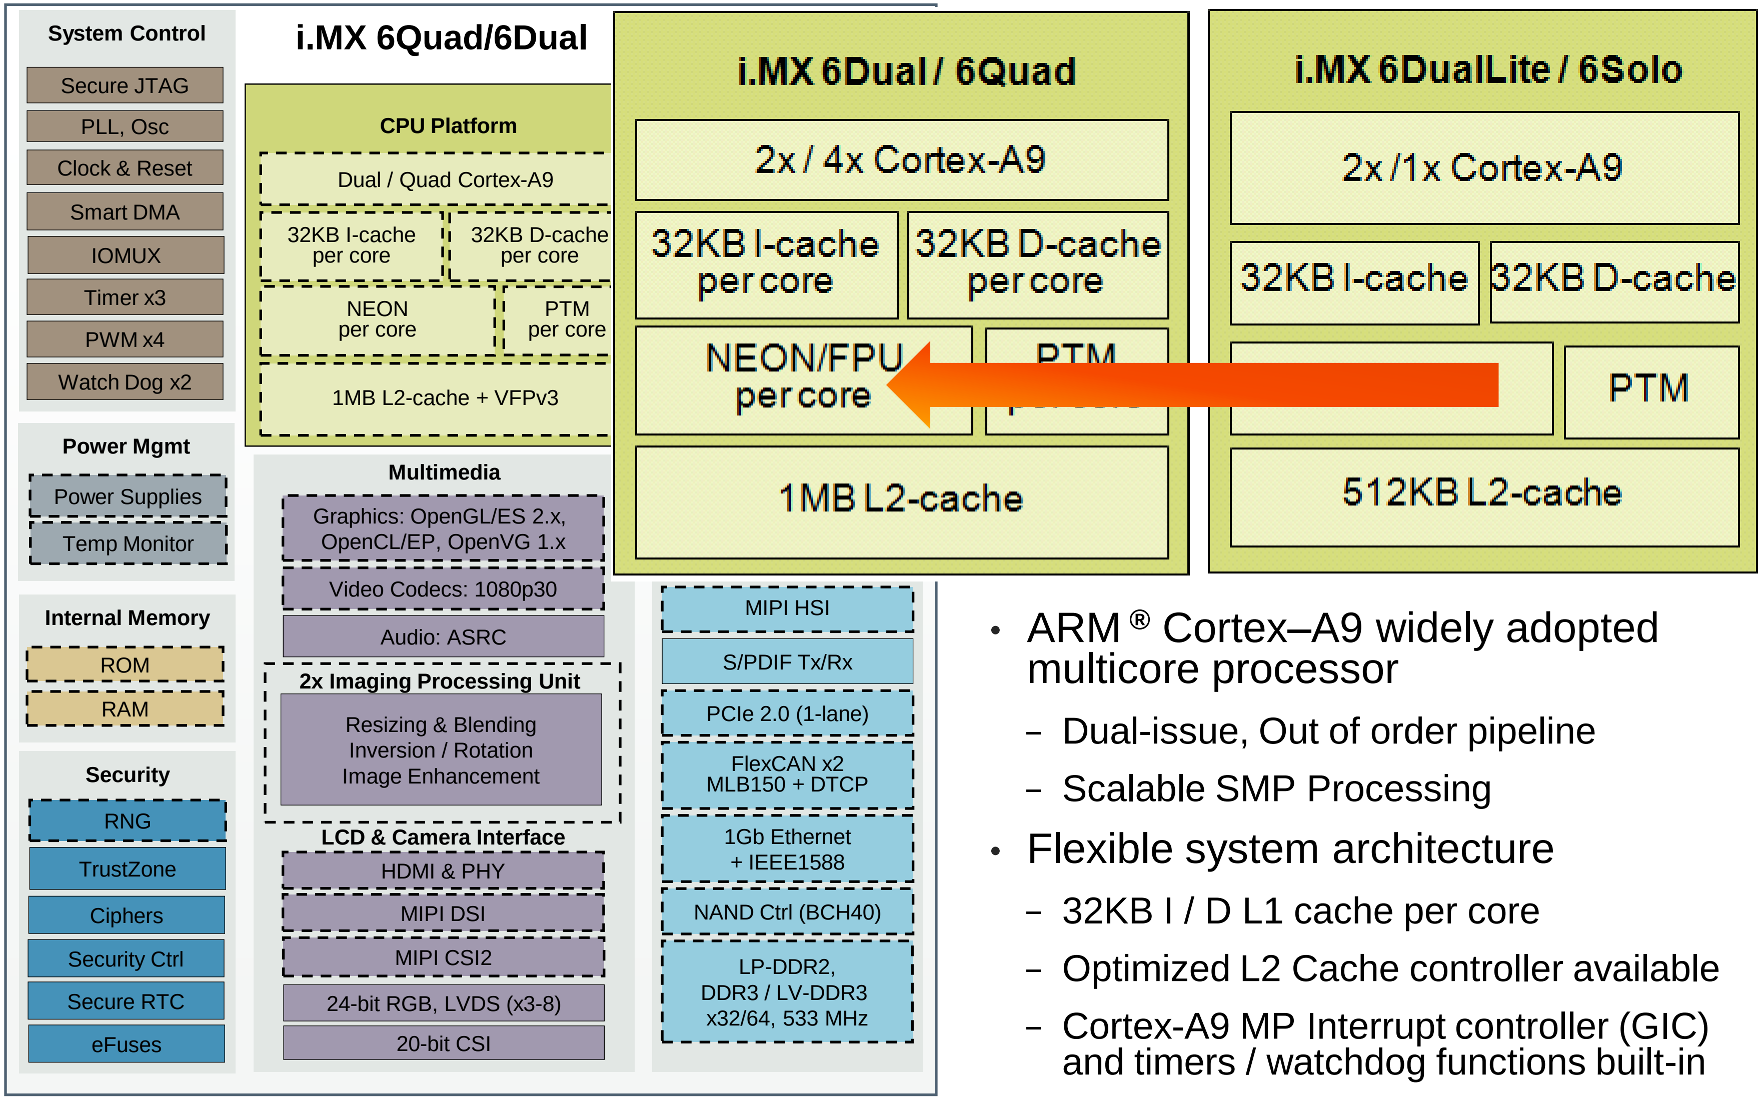

System Control

Secure JTAG

PLL, Osc

Clock & Reset

Smart DMA

IOMUX

Timer x3

PWM x4

Watch Dog x2

Power Mgmt

Power Supplies

Temp Monitor

Internal Memory

ROM

RAM

Security

RNG

TrustZone

Ciphers

Security Ctrl

Secure RTC

eFuses

i.MX 6Quad/6Dual

CPU Platform

Dual / Quad Cortex-A9

32KB I-cache per core

32KB D-cache per core

NEON per core

PTM per core

1MB L2-cache + VFPv3

Multimedia

Graphics: OpenGL/ES 2.x, OpenCL/EP, OpenVG 1.x

Video Codecs: 1080p30

Audio: ASRC

2x Imaging Processing Unit

Resizing & Blending
Inversion / Rotation
Image Enhancement

LCD & Camera Interface

HDMI & PHY

MIPI DSI

MIPI CSI2

24-bit RGB, LVDS (x3-8)

20-bit CSI

MIPI HSI

S/PDIF Tx/Rx

PCIe 2.0 (1-lane)

FlexCAN x2
MLB150 + DTCP

1Gb Ethernet
+ IEEE1588

NAND Ctrl (BCH40)

LP-DDR2,
DDR3 / LV-DDR3
x32/64, 533 MHz

i.MX 6Dual / 6Quad

2x / 4x Cortex-A9

32KB I-cache per core

32KB D-cache per core

NEON/FPU per core

PTM per core

1MB L2-cache

i.MX 6DualLite / 6Solo

2x /1x Cortex-A9

32KB I-cache

32KB D-cache

PTM

512KB L2-cache

- ARM® Cortex–A9 widely adopted multicore processor
  - Dual-issue, Out of order pipeline
  - Scalable SMP Processing
- Flexible system architecture
  - 32KB I / D L1 cache per core
  - Optimized L2 Cache controller available
  - Cortex-A9 MP Interrupt controller (GIC) and timers / watchdog functions built-in

ARM ® NEON Acceleration Unit

- NEON® is a wide SIMD (Single Instruction, Multiple Data) data processing architecture
  - Extension for ARM Cortex-A series processors and ARM instruction set
  - It provides support for the ARM v7 Advanced SIMD and Vector Floating-Point v3 (VFPv3) instruction sets.
- 32 registers, 64 bite wide (dual view as 16 registers, 128-bits wide)
- NEON instructions perform “Packed SIMD” processing
  - Registers are considered as vectors of elements of the same data type
  - Data types can be:
    - Singed/unsigned 8-bit, 16-bit, 32-bit 64-bit, single precision or floating point
    - Instructions perform same operation in all lane
- With the exception of simultaneous loads / stores, the processor can execute SIMD(VFP) instructions in parallel with ARM or Thumb instructions

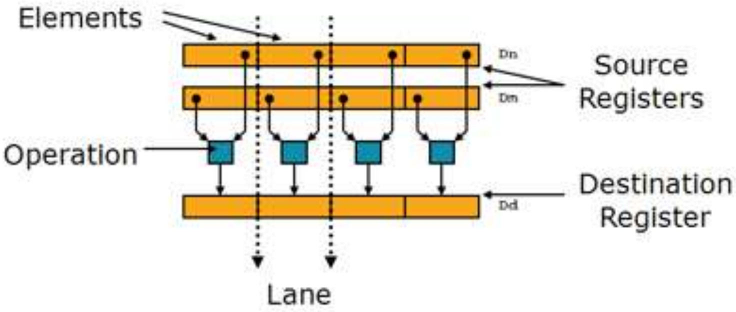

Software Enablement

- NEON can be enabled on Bare Metal Code
  - No OS required
  - Inline Assembler
  - NEON Intrinsics available (for GCC when mfpu=neon)
- Freescale offers Platform SDK
  - Bare metal SDK for i.Mx6 Series
  - starts Cortex-A9 (optional enable, NEON, L1 Cache, MMU, etc)
  - Reusable drivers and tests for many peripherals
  - https://community.freescale.com/docs/DOC-94139
- But how to enable NEON on an OS?

ARM ® Ne10 Project

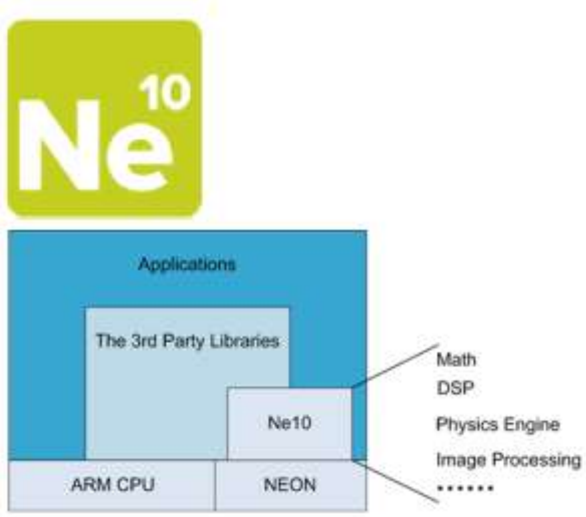

- Some facts:
  - An Open Optimized Software Library Project for the ARM architecture
  - The Ne10 library provides a set of the most commonly used functions that have been heavily optimized for ARM based CPUs with NEON
  - The library focuses on matrix/vector algebra and signal processing
  - Ne10 will evolve over time
  - License refer to http://projectne10.github.io/Ne10/doc/
- Software Enablement
  - Easy enablement on i.MX 6 series Linux or Android platforms
  - No assembly coding required to use libraries
  - Chose function from Ne10 pre-built library

Ne10 Supported Functions

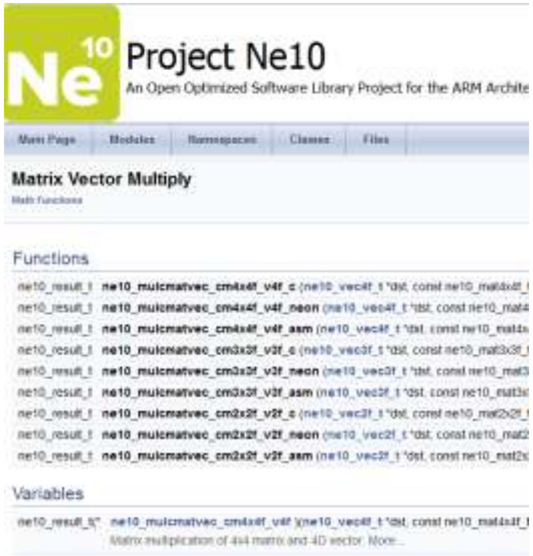

Project Ne10

An Open Optimized Software Library Project for the ARM Archite

Matrix Vector Multiply

Math functions

Functions

```
ne10_result_t  ne10_mulcmatvec_cm4x4f_v4f_c (ne10_vec4f_t *dst, const ne10_mat4x4f_t
ne10_result_t  ne10_mulcmatvec_cm4x4f_v4f_neon (ne10_vec4f_t *dst, const ne10_mat4
ne10_result_t  ne10_mulcmatvec_cm4x4f_v4f_asm (ne10_vec4f_t *dst, const ne10_mat4x
ne10_result_t  ne10_mulcmatvec_cm3x3f_v3f_c (ne10_vec3f_t *dst, const ne10_mat3x3f_t
ne10_result_t  ne10_mulcmatvec_cm3x3f_v3f_neon (ne10_vec3f_t *dst, const ne10_mat3
ne10_result_t  ne10_mulcmatvec_cm3x3f_v3f_asm (ne10_vec3f_t *dst, const ne10_mat3x
ne10_result_t  ne10_mulcmatvec_cm2x2f_v2f_c (ne10_vec2f_t *dst, const ne10_mat2x2f_t
ne10_result_t  ne10_mulcmatvec_cm2x2f_v2f_neon (ne10_vec2f_t *dst, const ne10_mat2
ne10_result_t  ne10_mulcmatvec_cm2x2f_v2f_asm (ne10_vec2f_t *dst, const ne10_mat2x
```

Variables

```
ne10_result_t(*  ne10_mulcmatvec_cm4x4f_v4f )(ne10_vec4f_t *dst, const ne10_mat4x4f_t
                 Matrix multiplication of 4x4 matrix and 4D vector. More...
```

Math Functions	Signal Processing Functions	Image Processing Functions
Vector Add / Sub	Complex FFT	Image Resize
Matrix Add / Sub	Float/Fixed point Complex FFT	Image Rotate
Vector Multiply / Multiply-Accumulator	Float/Fixed point Real2Complex FFT	
Matrix Multiply / Vector Mulitply	Finite Impulse Response (FIR) Filters	
Vector Div / Setc / Len /Normalize / Abs / Dot / Cross / Rsbc	Finite Impulse Response (FIR) Decimator	
Matrix Div / Determinant / Invertible / Transpose / Identify	Finite Impulse Response (FIR) Interpolator	
	Finite Impulse Response (FIR) Lattice Filters	
	Finite Impulse Response (FIR) Sparse Filters	
	Infinite Impulse Response (IIR) Lattice Filters	

Full documentation: http://projectne10.github.io/Ne10/doc/modules.html

Four Steps to Use Ne10 on i.MX 6 Series Running Linux

- Prepare target or host
  - ARM GCC, cmake, make,git
- Download Ne10 project
  - $ git clone git://github.com/projectNe10/Ne10
- Compile and use library
  - $cd $NE10PATH, mkdir build, cd build, cmake .., make
  - libNE10.a is placed in $NE10PATH /build/modules/
  - A test program "NE10_test_static" is placed in $NE10PATH /build/samples/.
  - More tests can be compiled $NE10PATH /build/test/
- You can run it

Demo on i.MX 6Quad Sabre SDB

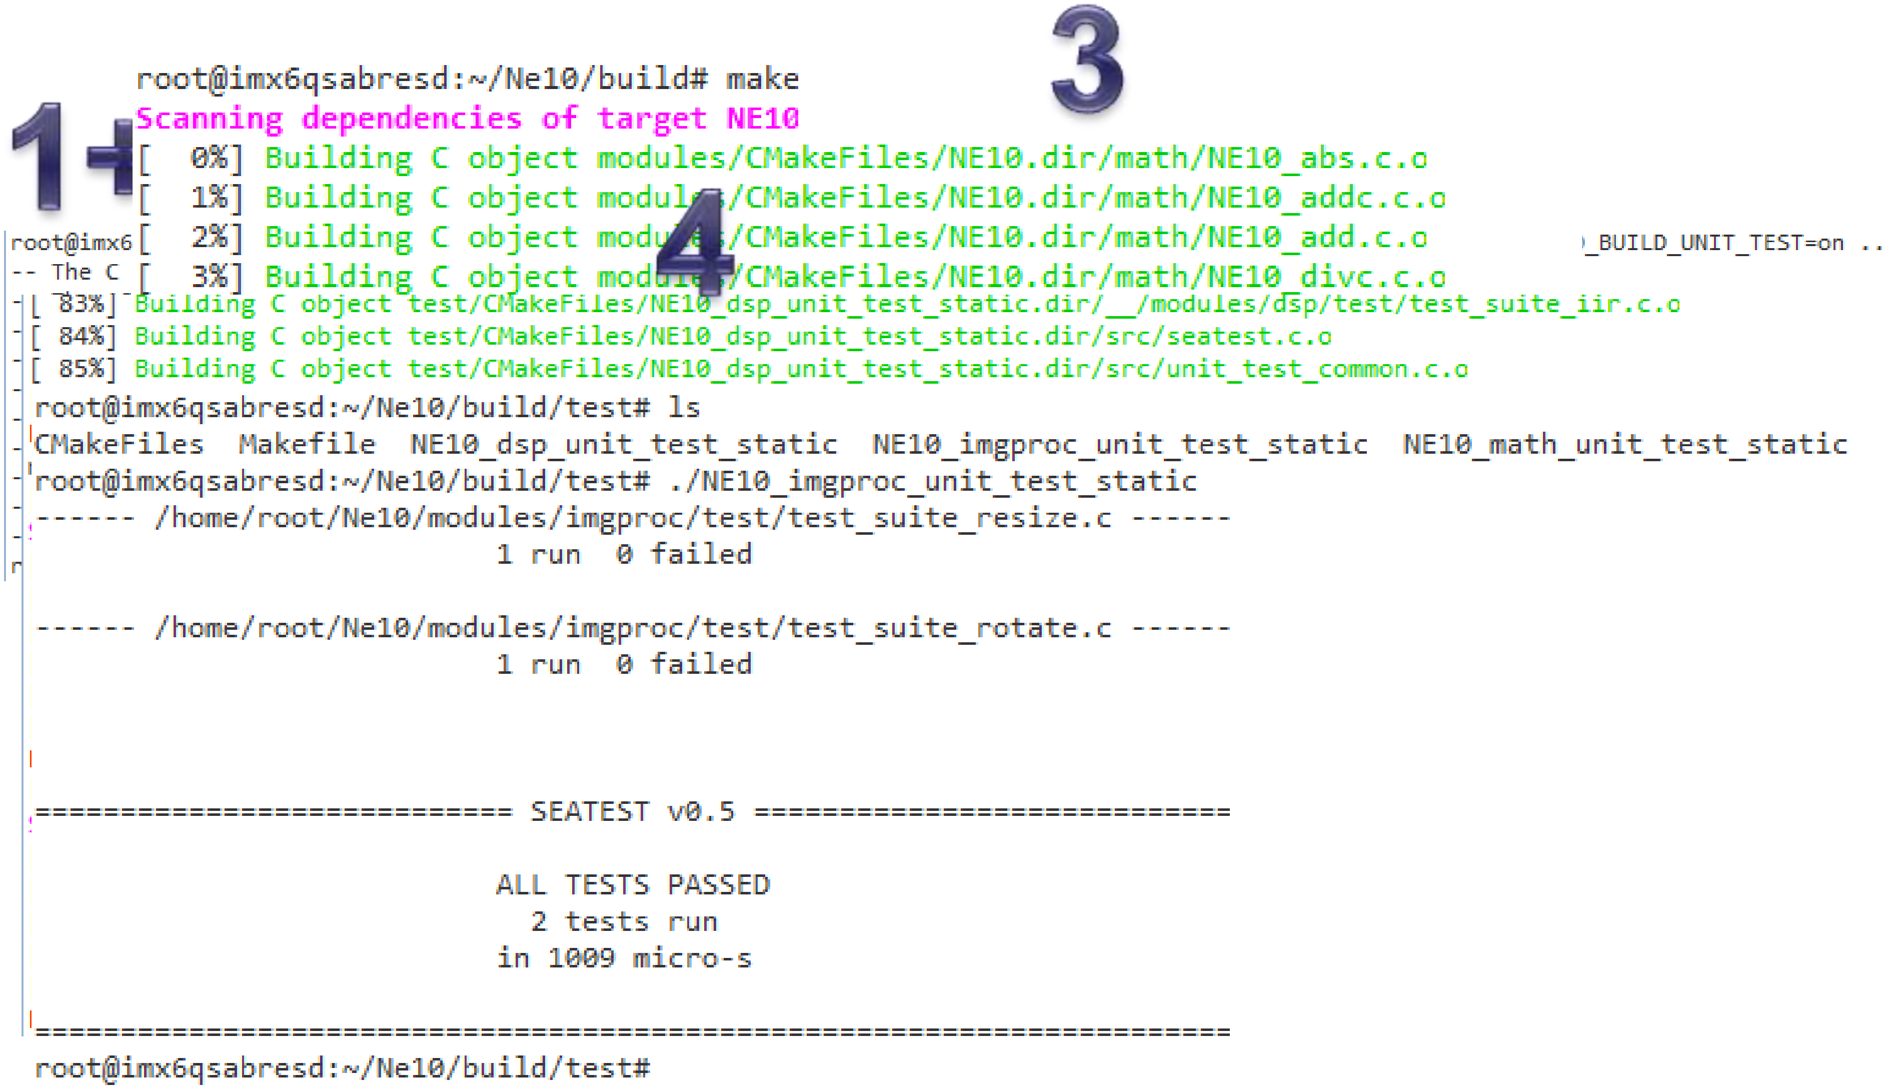

```
root@imx6qsabresd:~/Ne10/build# make
Scanning dependencies of target NE10
[  0%] Building C object modules/CMakeFiles/NE10.dir/math/NE10_abs.c.o
[  1%] Building C object modules/CMakeFiles/NE10.dir/math/NE10_addc.c.o
[  2%] Building C object modules/CMakeFiles/NE10.dir/math/NE10_add.c.o
[  3%] Building C object modules/CMakeFiles/NE10.dir/math/NE10_divc.c.o
```

```
root@imx6                                                    _BUILD_UNIT_TEST=on ..
-- The C
```

```
[ 83%] Building C object test/CMakeFiles/NE10_dsp_unit_test_static.dir/__/modules/dsp/test/test_suite_iir.c.o
[ 84%] Building C object test/CMakeFiles/NE10_dsp_unit_test_static.dir/src/seatest.c.o
[ 85%] Building C object test/CMakeFiles/NE10_dsp_unit_test_static.dir/src/unit_test_common.c.o
```

```
root@imx6qsabresd:~/Ne10/build/test# ls
CMakeFiles  Makefile  NE10_dsp_unit_test_static  NE10_imgproc_unit_test_static  NE10_math_unit_test_static
root@imx6qsabresd:~/Ne10/build/test# ./NE10_imgproc_unit_test_static
------ /home/root/Ne10/modules/imgproc/test/test_suite_resize.c ------
                           1 run  0 failed

------ /home/root/Ne10/modules/imgproc/test/test_suite_rotate.c ------
                           1 run  0 failed




============================ SEATEST v0.5 ============================

                           ALL TESTS PASSED
                             2 tests run
                           in 1009 micro-s


======================================================================
root@imx6qsabresd:~/Ne10/build/test#
```

i.MX 6Quad/6Dual Applications Processor

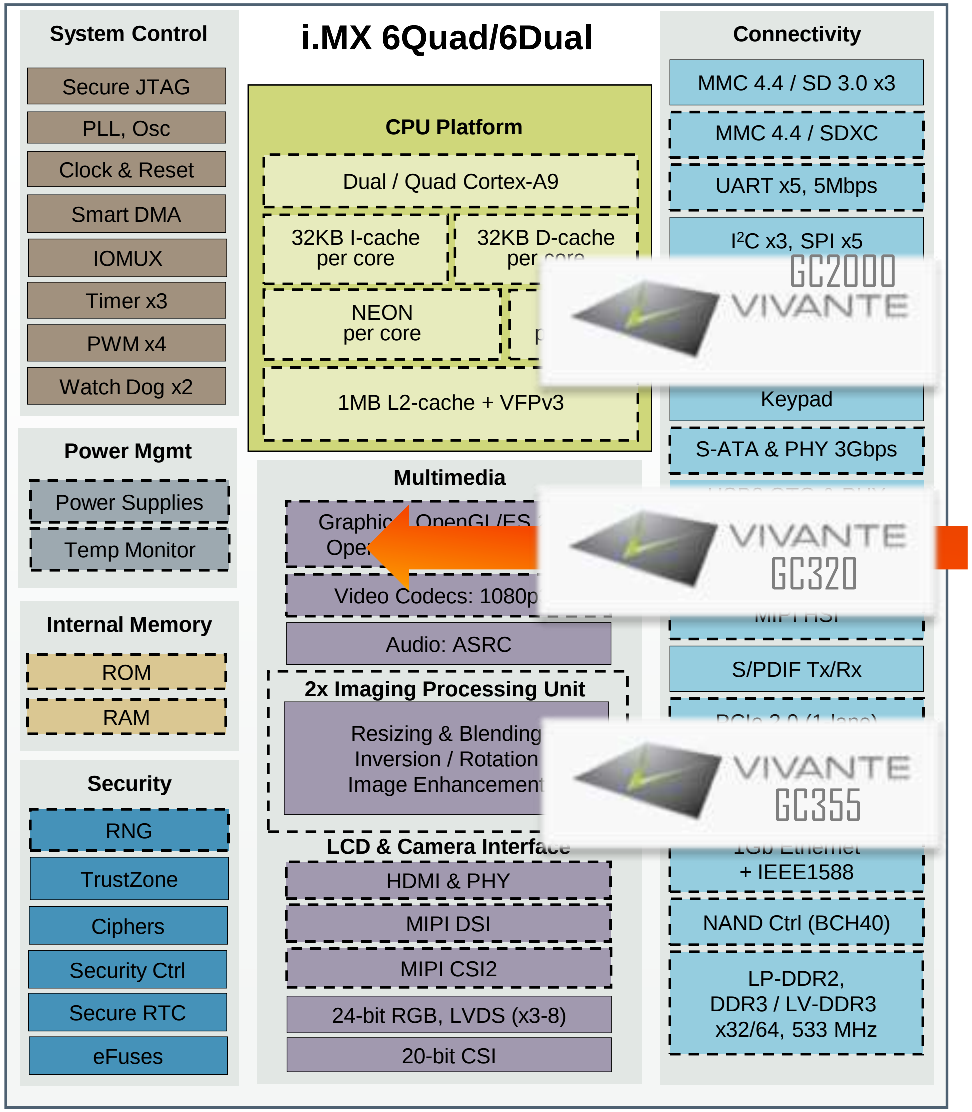

i.MX 6Quad/6Dual

System Control
Secure JTAG
PLL, Osc
Clock & Reset
Smart DMA
IOMUX
Timer x3
PWM x4
Watch Dog x2

Power Mgmt
Power Supplies
Temp Monitor

Internal Memory
ROM
RAM

Security
RNG
TrustZone
Ciphers
Security Ctrl
Secure RTC
eFuses

CPU Platform
Dual / Quad Cortex-A9
32KB I-cache per core
32KB D-cache per core
NEON per core
p
1MB L2-cache + VFPv3

Multimedia
Graphic OpenGL/ES
Ope
Video Codecs: 1080p
Audio: ASRC
2x Imaging Processing Unit
Resizing & Blending
Inversion / Rotation
Image Enhancement
LCD & Camera Interface
HDMI & PHY
MIPI DSI
MIPI CSI2
24-bit RGB, LVDS (x3-8)
20-bit CSI

Connectivity
MMC 4.4 / SD 3.0 x3
MMC 4.4 / SDXC
UART x5, 5Mbps
I2C x3, SPI x5
Keypad
S-ATA & PHY 3Gbps
MIPI HSI
S/PDIF Tx/Rx
1Gb Ethernet + IEEE1588
NAND Ctrl (BCH40)
LP-DDR2, DDR3 / LV-DDR3 x32/64, 533 MHz

GC2000 VIVANTE
VIVANTE GC320
VIVANTE GC355

- Introducing the i.MX Triple-Play Graphics Solution
- High-Performance OpenGL/OpenCL GPU Core
  - Vivante GC2000 @ 528MHz (600 MHz shader)
  - 88M triangles / sec (176MT/s @ 50% cull)
  - 4 Shader Cores: 21.6 GFLOPS
  - 1056 M Pixels/s, 1056 M Texels/ s
- High-Performance Composition Engine 2D Core
  - Up to 633M pixels / sec raw performance
  - Vivante GC320
- High-Performance Vector Graphics Core
  - 264M Pixels / sec raw performance
  - Native OpenVG 1.1 vector acceleration with hardware tessellation
  - Vivante GC355

Vivante GC2000 ScalarMorphic™ Architecture

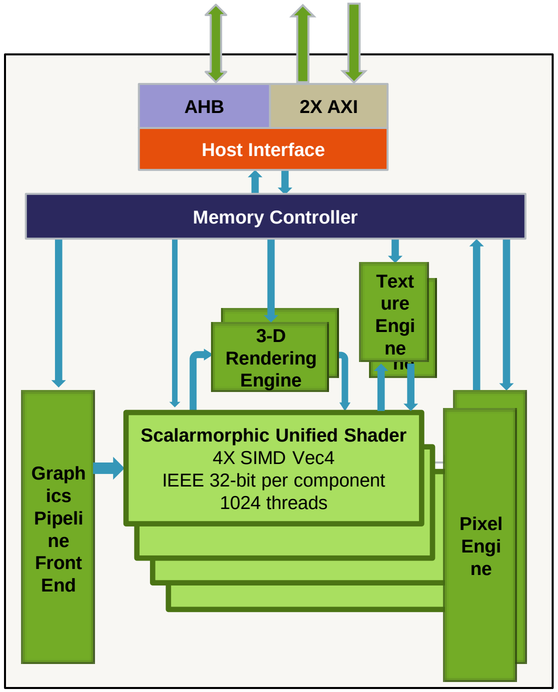

- Architectural Features
- CoolThreads™ multilevel clock gating and massive data parallelism
- 4X SIMD Vec-4 unified shaders (IEEE single precision FP)
- Point sample, bi-linear/tri-linear and cubic/volumetric textures
- Balanced performance/ bandwidth
- Memory locality / Texture compression
- Frame buffer/Z compression
- Native OpenGL ES 2.0 rendering
- Quality 4x MSAA anti-aliasing
- Depth buffering

- System Functionality
- Up to 1024 independent threads
- Low bandwidth characteristics for both simple and complex scenes
- Comprehensive API support
- OpenGL ES 1.1 / 2.0
- OpenVG 1.1
- OpenGL 2.1 & 3.0 (w/ GLX)
- OpenCL 1.1 EP

i.MX 6 Series GPU Driver Overview

- The GPU driver is structured in a “User space” part and a “Kernel space” part that communicate through ioctls
- The user side contains command buffers where the application puts commands for the GPU
- On the kernel side, command queues are used to link and manage the command buffers supplied
- The hardware is synchronized with the software through hardware events

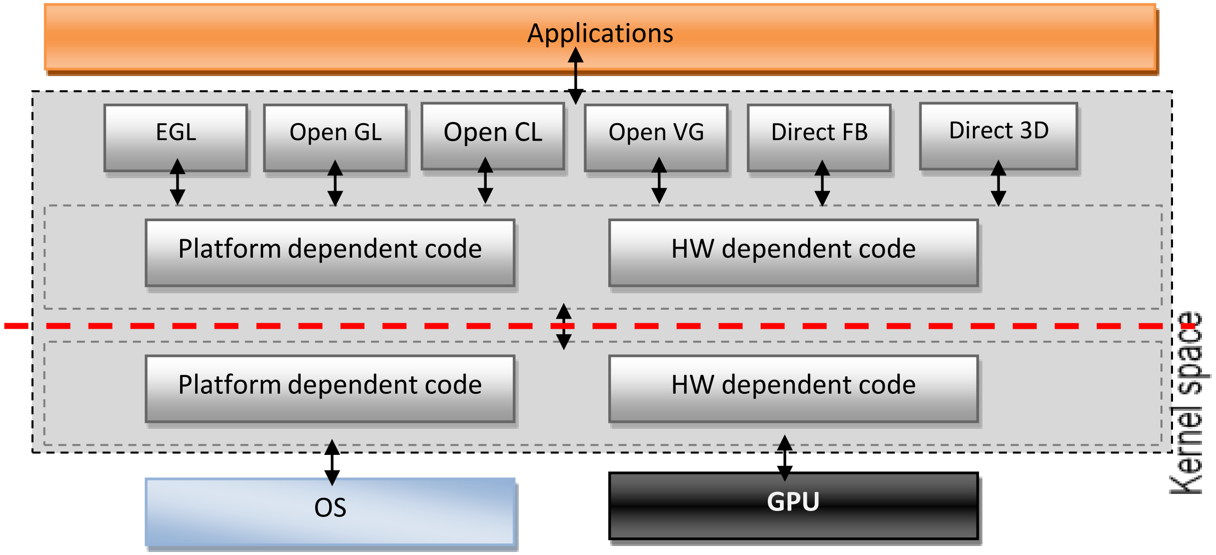

OpenCL Overview

- OpenCL is an open specification API from the Khronos Group
  - Same standards body that made OpenGL, OpenVG, OpenMAX, etc.
- OpenCL enables asynchronous multi-core programming for heterogeneous computing environments
  - Allows the user to run work ‘kernels’ that run over large data sets and across many processors
  - Effectively makes GPU look like a cluster of DSPs executing on a task
  - i.MX 6Q/D implements the OpenCL 1.1 Embedded Profile
- Language Specification
  - C-based cross-platform programming interface
  - Subset of ISO C99 with language extensions - familiar to developers
  - Well-defined numerical accuracy (IEEE 754 rounding with specified max error)
  - Online or offline compilation and build of compute kernel executables
  - Includes a rich set of built-in functions
- Platform Layer API
  - A hardware abstraction layer over diverse computational resources
  - Query, select and initialize compute devices
  - Create compute contexts and work-queues
- Runtime API
  - Execute compute kernels
  - Manage scheduling, compute, and memory resources

OpenCL Example Use Cases

OpenCL efficient functions

- Haar-like Features / Wavelets
- Sobel Line Detection
- Computing Eigenvalues
- FFT
- Noise Generation
- DCT
- Histogram Generation
- Histogram of Oriented Gradients (HOGs)
- Demosiac

i.MX 6 Series Open CL Execution Model

Compute Units per device (Shader cores)	**4**
Processing Elements per compute unit	4
Profile	Embedded
Preferred work-group/ thread group size	16
Max count global work-items each dim	64K
Max count of work-items each dim per work-group	1K
Instruction Memory	512
Texture Samplers available to OCL	4
L1 Cache Size	4KB

Power Versus Performance Tradeoff

- Simple test on i.MX 6Quad SabreSDB – Yocto 3.10.17 Beta Release
- Enabled two Cortex-A9 cores
- Vector ADD (Data type float32)
  - Cortex-A9
  - Cortex-A9 + NEON
  - Cortex-A9 + OpenCL (GC2000)
- Depending on size NEON / OpenCL can show significant performance boost
- On small data size OpenCL sub optimal or even worse because of OpenCL overhead
- Significant power consumption increase (~+60%) when enabling OpenCL

Best Reasons When to Chose NEON / GPU OpenCL

- When CPU utilization is high due to DSP and MATH computing
  - Filter (IIR, FIR) MAC, Floats
- When dealing with vectors
- NEON runs at CPU speed and takes advantage of L1/L2 cache
- CPU and NEON can make use of local SRAM
- Make use of Ne10 Library to use NEON for Linux or Android applications
- OpenCL can bring more performance and is more flexible

Connect Peripherals to i.MX 6

Why Attach or Integrate Peripherals in Industrial?

- Industrial IO Systems
  - Support (different) industrial Ethernet protocols
  - Require often Integration of proprietary ASIC or FPGA
  - Flexible and efficient interface
- Compact form factors
  - CPU needs to service additional channels (UART, USB, etc)
- Cooling due to compact size or closed body an issue
  - Power efficient system architectures required
- At the same time data throughput demand growing
  - An high throughput interface is required

i.MX 6 Series Interface Offering

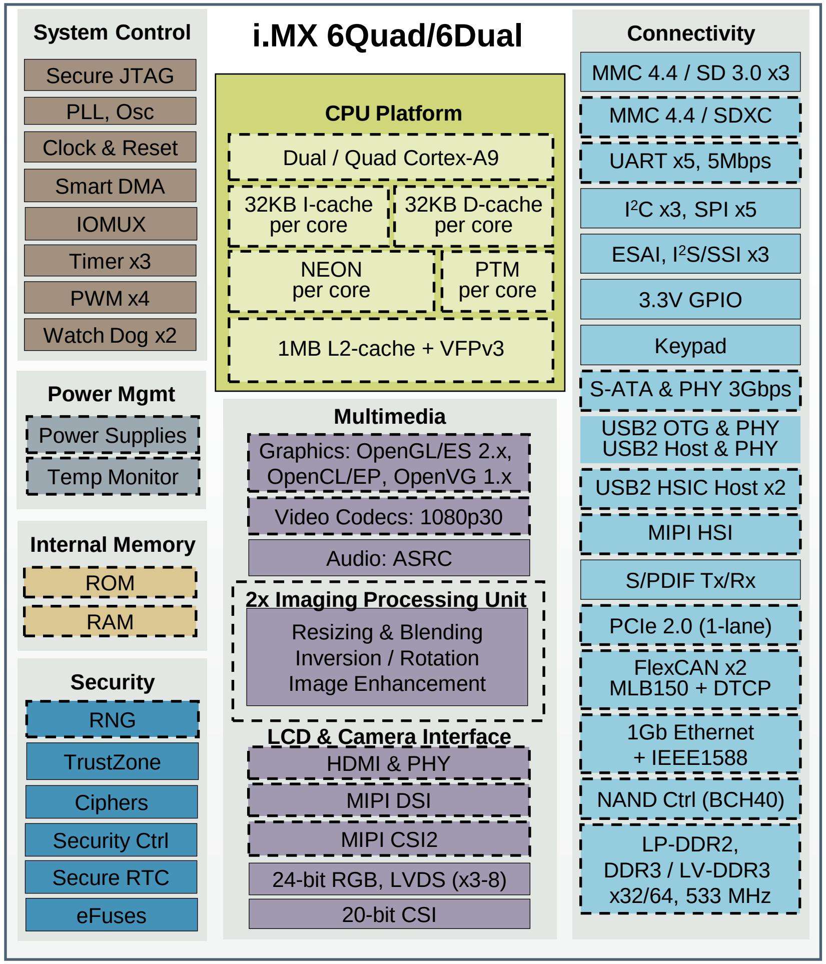

i.MX 6Quad/6Dual

System Control

Secure JTAG

PLL, Osc

Clock & Reset

Smart DMA

IOMUX

Timer x3

PWM x4

Watch Dog x2

Power Mgmt

Power Supplies

Temp Monitor

Internal Memory

ROM

RAM

Security

RNG

TrustZone

Ciphers

Security Ctrl

Secure RTC

eFuses

CPU Platform

Dual / Quad Cortex-A9

32KB I-cache per core

32KB D-cache per core

NEON per core

PTM per core

1MB L2-cache + VFPv3

Multimedia

Graphics: OpenGL/ES 2.x, OpenCL/EP, OpenVG 1.x

Video Codecs: 1080p30

Audio: ASRC

2x Imaging Processing Unit

Resizing & Blending

Inversion / Rotation

Image Enhancement

LCD & Camera Interface

HDMI & PHY

MIPI DSI

MIPI CSI2

24-bit RGB, LVDS (x3-8)

20-bit CSI

Connectivity

MMC 4.4 / SD 3.0 x3

MMC 4.4 / SDXC

UART x5, 5Mbps

I^2C x3, SPI x5

ESAI, I^2S/SSI x3

3.3V GPIO

Keypad

S-ATA & PHY 3Gbps

USB2 OTG & PHY

USB2 Host & PHY

USB2 HSIC Host x2

MIPI HSI

S/PDIF Tx/Rx

PCIe 2.0 (1-lane)

FlexCAN x2

MLB150 + DTCP

1Gb Ethernet

+ IEEE1588

NAND Ctrl (BCH40)

LP-DDR2,

DDR3 / LV-DDR3

x32/64, 533 MHz

- i.MX 6 series Industrial Ethernet support (partners)
  - EtherCAT master
  - PROFINET RT (Not IRT)
  - Modbus TCP
  - IEEE 1588
  - DeviceNet (CAN)
- EtherCAT slave or other proprietary Interfaces require ASIC / FPGA

External Interface Module (EIM)

- The EIM handles the interface to devices peripherals and memory to the chip
- Feature Highlights:
  - Up to four chip selects for external devices
  - Support for multiplexed address / data bus operation
  - Programmable wait-state generator for each CS
  - Asynchronous or Synchronous access to devices
  - Flexible 8/16 or 32-bit configurable data width
  - Asynch. page Mode access and Burst support for synch mode
  - External Interrupt Support

EIM Benefit and Software Enablement

- Connect FPGA / ASIC directly through i.MX 6 series EIM
- No protocol overhead thus efficient when accessing small data junks (single bytes, half or words)
- If real-time data are computed in external peripheral they can be transferred with low latency to i.MX 6 series (could run Linux RT Xenmai)
- To burst data is are possible
- Driver can directly memory map EIM address space
- Freescale provides reference driver in Platform SDK
  - Bare metal driver can be ported to Linux, WinCE or RTOS

i.MX 6 Series Interface Offering

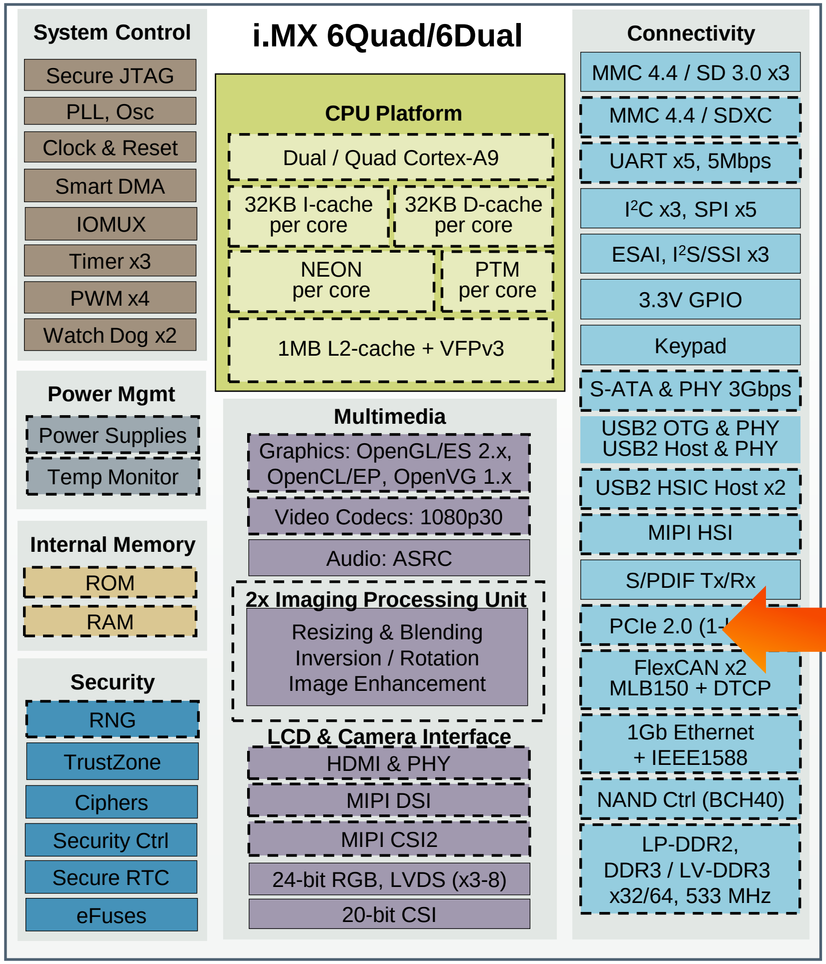

i.MX 6Quad/6Dual

System Control
- Secure JTAG
- PLL, Osc
- Clock & Reset
- Smart DMA
- IOMUX
- Timer x3
- PWM x4
- Watch Dog x2

Power Mgmt
- Power Supplies
- Temp Monitor

Internal Memory
- ROM
- RAM

Security
- RNG
- TrustZone
- Ciphers
- Security Ctrl
- Secure RTC
- eFuses

CPU Platform
- Dual / Quad Cortex-A9
- 32KB I-cache per core
- 32KB D-cache per core
- NEON per core
- PTM per core
- 1MB L2-cache + VFPv3

Multimedia
- Graphics: OpenGL/ES 2.x, OpenCL/EP, OpenVG 1.x
- Video Codecs: 1080p30
- Audio: ASRC

2x Imaging Processing Unit
- Resizing & Blending
- Inversion / Rotation
- Image Enhancement

LCD & Camera Interface
- HDMI & PHY
- MIPI DSI
- MIPI CSI2
- 24-bit RGB, LVDS (x3-8)
- 20-bit CSI

Connectivity
- MMC 4.4 / SD 3.0 x3
- MMC 4.4 / SDXC
- UART x5, 5Mbps
- I²C x3, SPI x5
- ESAI, I²S/SSI x3
- 3.3V GPIO
- Keypad
- S-ATA & PHY 3Gbps
- USB2 OTG & PHY
- USB2 Host & PHY
- USB2 HSIC Host x2
- MIPI HSI
- S/PDIF Tx/Rx
- PCIe 2.0 (1-
- FlexCAN x2
- MLB150 + DTCP
- 1Gb Ethernet + IEEE1588
- NAND Ctrl (BCH40)
- LP-DDR2, DDR3 / LV-DDR3 x32/64, 533 MHz

- But what if higher throughput is required?

PCI EXPRESS™

PCI Express Overview

- PCIe Specification
  - http://www.pcisig.com/specifications/pciexpress/
- PCI Express architecture is a high performance, IO interconnect for peripherals in computing/communication platforms
- Evolved from PCI and PCI-X architectures
  - 32-bit or 64-bit data bus, addressing scheme
  - Bus arbitration
- PCIe enhancements
  - Point-to-point connection
  - Serial high Speed Interface (1,2,4,8,16, etc Lanes)
  - PCIe uses differenc encoding shemes like 8b/10b
  - Higher throughput
- PCIe Topology
  - PCIe Root Complex (typical Master or main CPU)
  - PCIe Endpoint (typical high speed Peripheral)
  - PCIe Switch (Peripheral Expander)

i.MX 6 Series PCIe Integration

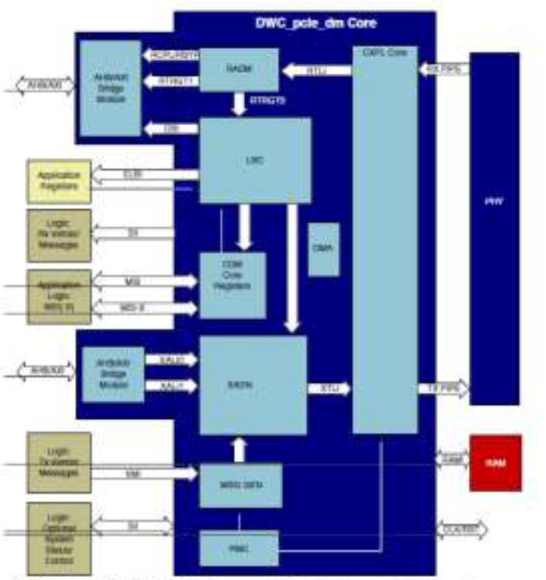

Note: Optional internal (iATU) and external (eATU) Address Translation Units/Interfaces not shown.

Note: Narrow arrows represent response signal paths to requests (broad arrows)

Note: In RC Mode the ELBI pins are present but not operational.

Figure 48-1. DM Core Block Diagram (with AXI Bridge Module)

- PCIe include the following cores
  - PCIe Dual Mode (DM) core
  - PCIe Root Complex (RC) core
  - PCIe Endpoint (EP) core
  - AXI Bridge to Interconnect

- PCI Express PHY supports
  - PCIe 2.0 / 5Gbps
  - PCIe 1.1 / 2.5 Gbps
  - IEEE 1149.6 (JTAG) Bscan
  - Built in self Tester (BIST)

PCIe Protocol Stack

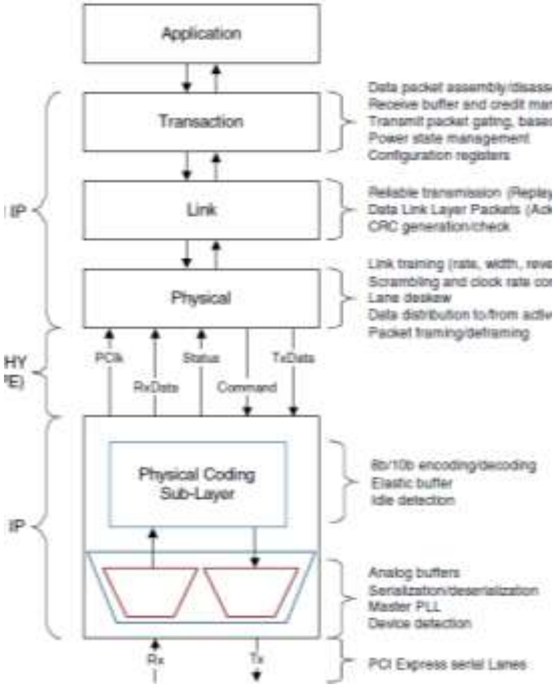

- PCI Express Dual Mode controller and PCIe PHY hide most of the complexity of the protocol from application
- PCIe complex implements standard PIPE interface between digital logic (PCIe controller) and interface PHY
- i.MX 6 series BSP supports PCIe

How to Enable PCIe Driver in FSL Linux

- i.MX 6 series FSL Yocto 3.10.17 release supports PCIe
  - Root Complex or Endpoint Driver

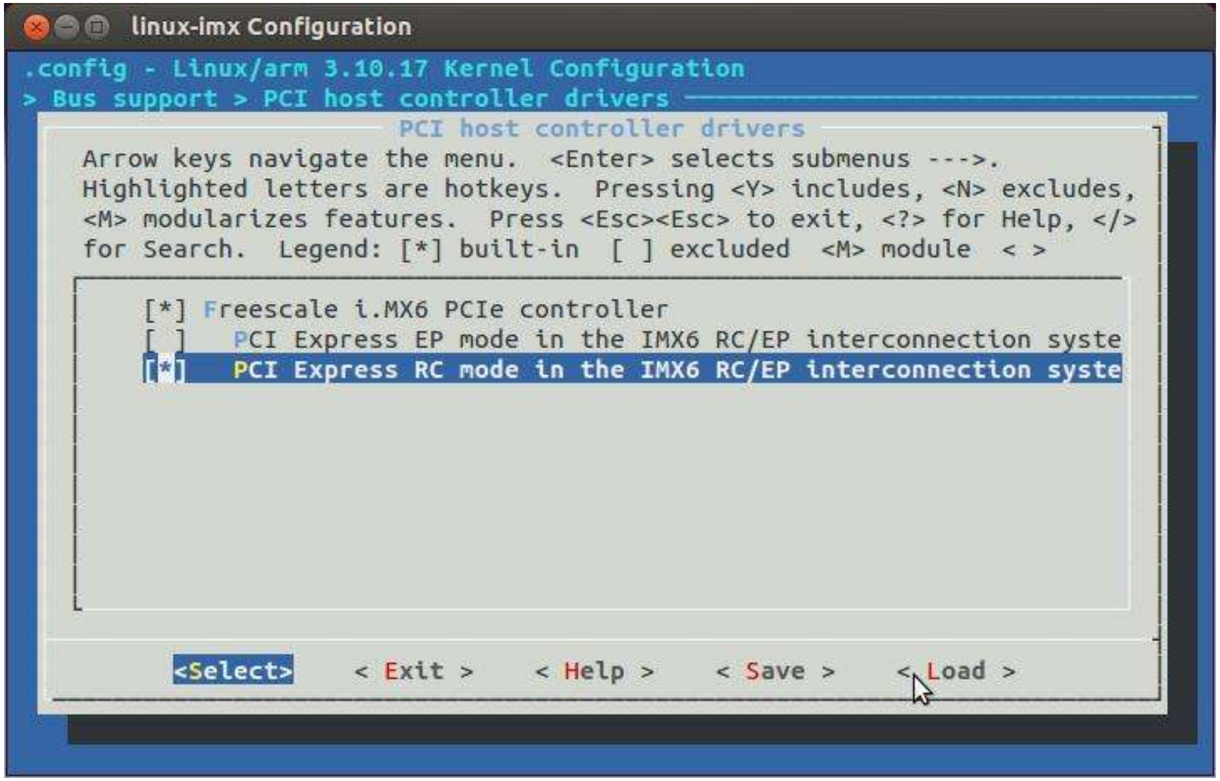

```
.config - Linux/arm 3.10.17 Kernel Configuration
> Bus support > PCI host controller drivers
                    PCI host controller drivers
  Arrow keys navigate the menu.  <Enter> selects submenus --->.
  Highlighted letters are hotkeys.  Pressing <Y> includes, <N> excludes,
  <M> modularizes features.  Press <Esc><Esc> to exit, <?> for Help, </>
  for Search.  Legend: [*] built-in  [ ] excluded  <M> module  < >

      [*] Freescale i.MX6 PCIe controller
      [ ]   PCI Express EP mode in the IMX6 RC/EP interconnection syste
      [*]   PCI Express RC mode in the IMX6 RC/EP interconnection syste
```

Demo of Using and Testing PCIe WLAN Card

- Follow Linux RM – PCIe How to test
  - Generic IEEE 802.11 Networking Stack (mac80211) used by WIFI devices
    - Symbol: MAC80211 [=y] Type : tristate Prompt: Generic IEEE 802.11 Networking Stack (mac80211) Defined at net/mac80211/Kconfig:1 Depends on: NET [=y] && WIRELESS [=y] && CFG80211 [=y] Location: -> Networking support (NET [=y]) -> Wireless (WIRELESS [=y])
    - CONFIG_IWLAGN: Select to build the driver supporting the: Intel Wireless Wi-Fi Link Next-Gen AGN
    - WIFI firmware configurations:
      - iwlagn: Intel(R) Wireless WiFi Link AGN driver for Linux, in-tree:
      - iwlagn: Copyright(c) 2003-2011 Intel Corporation
      - PCI: enabling device 0000:01:00.0 (0140 -> 0142)
      - iwlagn 0000:01:00.0: Detected Intel(R) Centrino(R) Ultimate-N 6300 AGN, REV=0x74
      - iwlagn 0000:01:00.0: device EEPROM VER=0x43a, CALIB=0x6
      - iwlagn 0000:01:00.0: Device SKU: 0Xb
      - iwlagn 0000:01:00.0: Valid Tx ant: 0X7, Valid Rx ant: 0X7
      - iwlagn 0000:01:00.0: Tunable channels: 13 802.11bg, 24 802.11a channels

i.MX 6 Series PCI Express Throughput Results

	ARM Core used as the bus master and cache disabled	ARM Core used as the bus master and cache enabled	IPU used as the bus master (DMA)
Data size in one write tlp	8 bytes	32 bytes	64 bytes
Write speed	~109 MB/s	~298 MB/s	~344 MB/s
Data size in one read tlp	32 bytes	64 bytes	64 bytes
Read speed	~29 MB/s	~100 MB/s	~211 MB/s

Summary

- i.MX 6 series is an successful and optimal choice for industrial applications
  - i.MX 6Quad to i.MX 6Solo full pin to pin compatible
    - Industrial qualified devices
    - Part of the Freescale product longevity program
  - i.MX 6 series offers several computing engines
    - ARM Cortex-A9 including NEON / VFPv3
    - Vivante GP GPU supporting OpenCL
  - i.MX 6 series rich and flexible interfaces
    - Choose between flexible and low latency (EIM) vs. high data throughput (PCIe)

freescale™

www.Freescale.com