

Image Processing Applications for Heterogeneous Computing Architectures

Dietmar Fey

Friedrich-Alexander-Universität
Erlangen-Nürnberg (FAU)

KeyNote, SIGNAL 2016, Lisbon,
2016 June 30



FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG

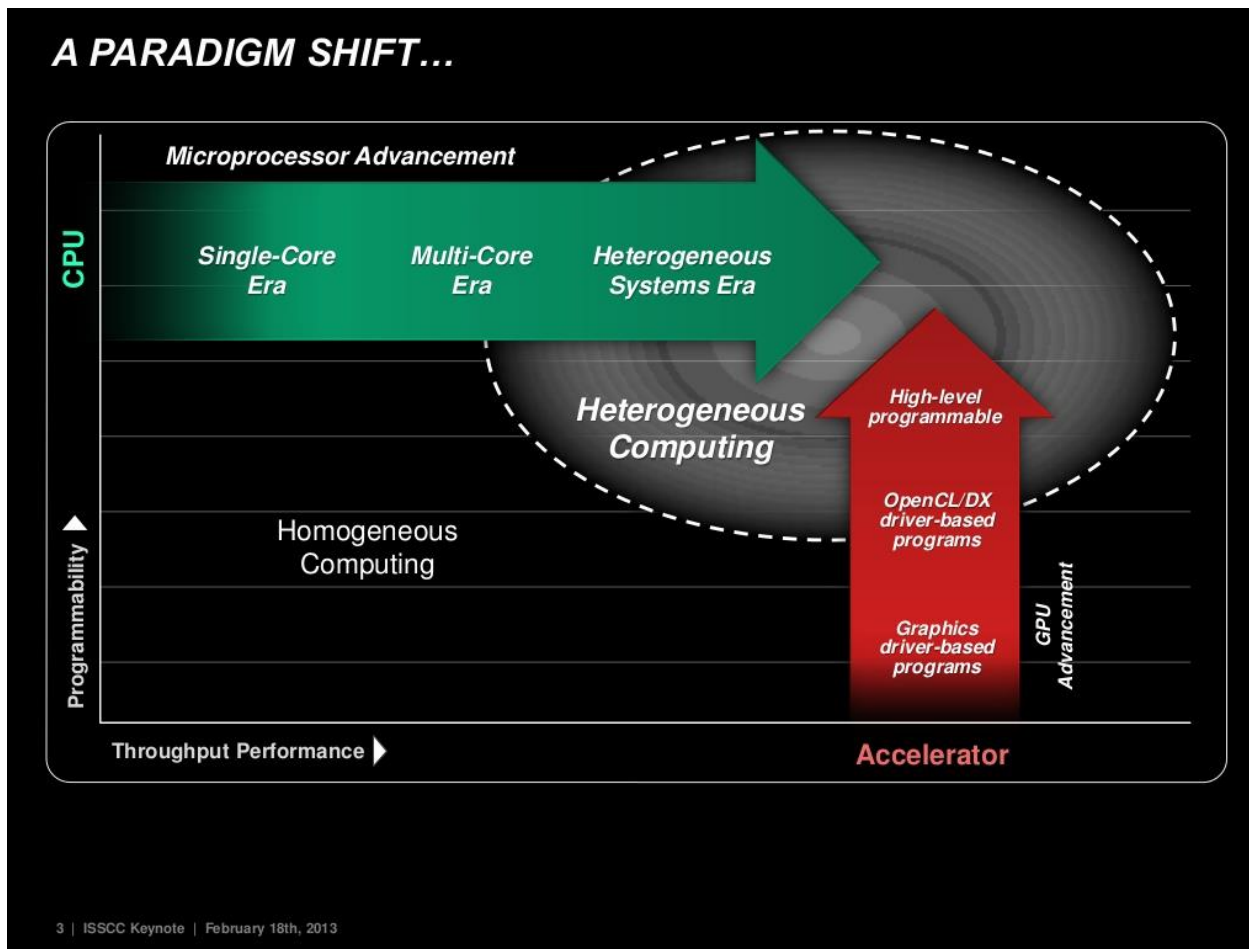
TECHNISCHE FAKULTÄT

Outline

1. The „Evolution“ to heterogeneous computing
2. Heterogeneous computing in embedded systems
3. Examples and challenges for embedded HIS
 - Near-term: Smart Cameras
 - More ambitious – Mid-term / long-term: Sensor fusion in automotive
4. Challenges on software and design side for heterogeneous computing
 - Heterogeneous design platforms
 - Heterogeneous programming environments

The „Evolution“ to heterogeneous computing?

Statement of the Heterogeneous Systems Foundation member

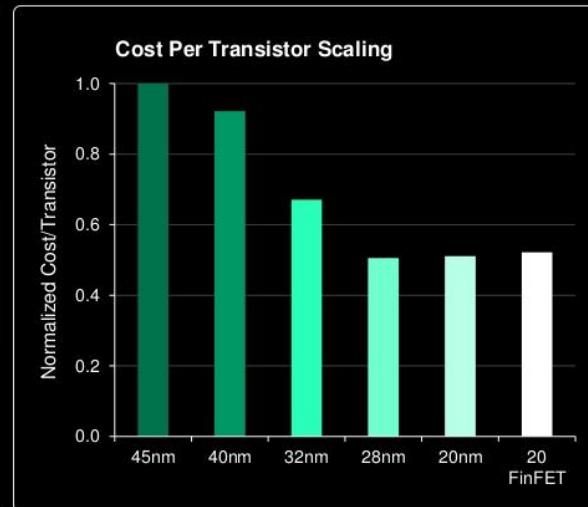
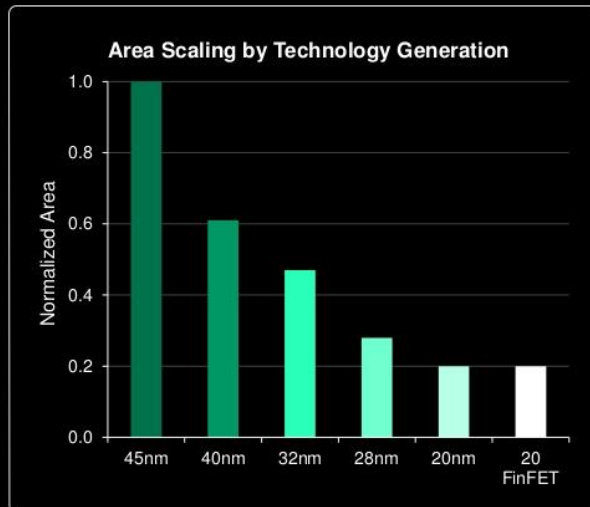


Slide from Dr. Lisa Su,
Vice President AMD

The „Evolution“ to heterogeneous computing?

Statement of the Heterogeneous Systems Foundation member

CHALLENGES TO MOORE'S LAW SCALING




- Lithography challenges begin severely limiting area scaling at 20nm node
 - Fewer 1X metals due to cost
 - Less aggressive feature scaling due to lithography challenges
- Compounded by rapidly increasing lithography costs
 - 28 → 20nm transition is inflection point with dual exposure
 - No cost / transistor crossover for first time at 28 → 20nm transition

Slide from Dr. Lisa Su,
Vice President AMD

The „Evolution“ to heterogeneous computing?

Statement of the Heterogeneous Systems Foundation member



Slide from Dr. Lisa Su,
Vice President AMD

The „Evolution“ to heterogeneous computing?

Three main reasons to my view

Energy

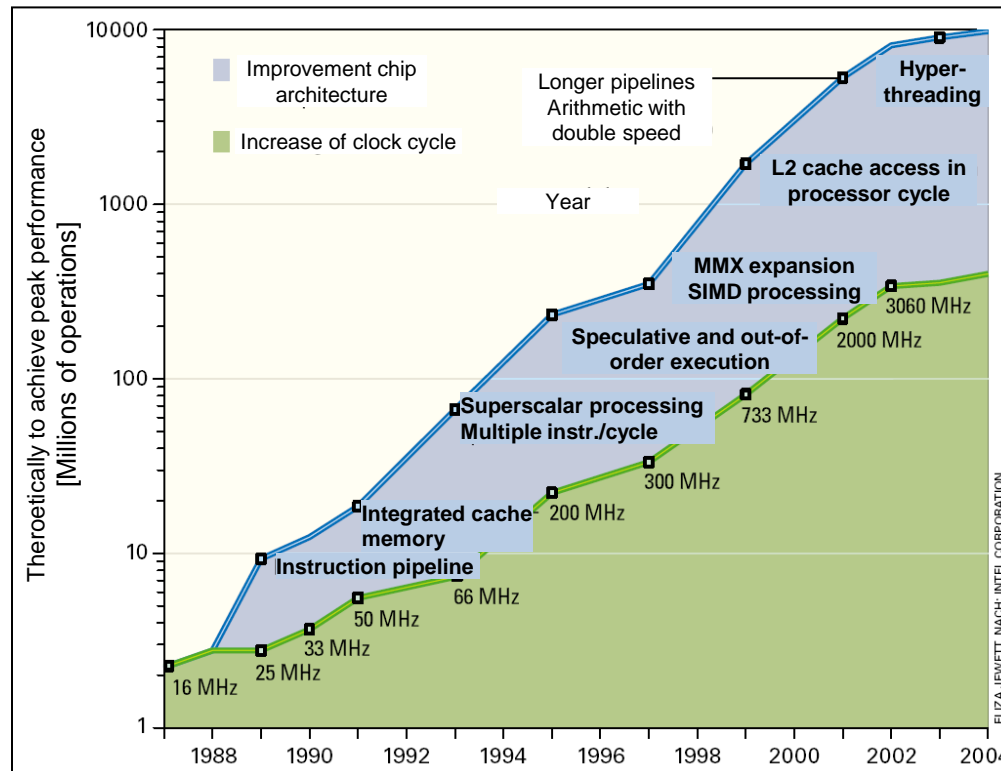
Energy

Energy

The „Evolution“ to heterogeneous computing?

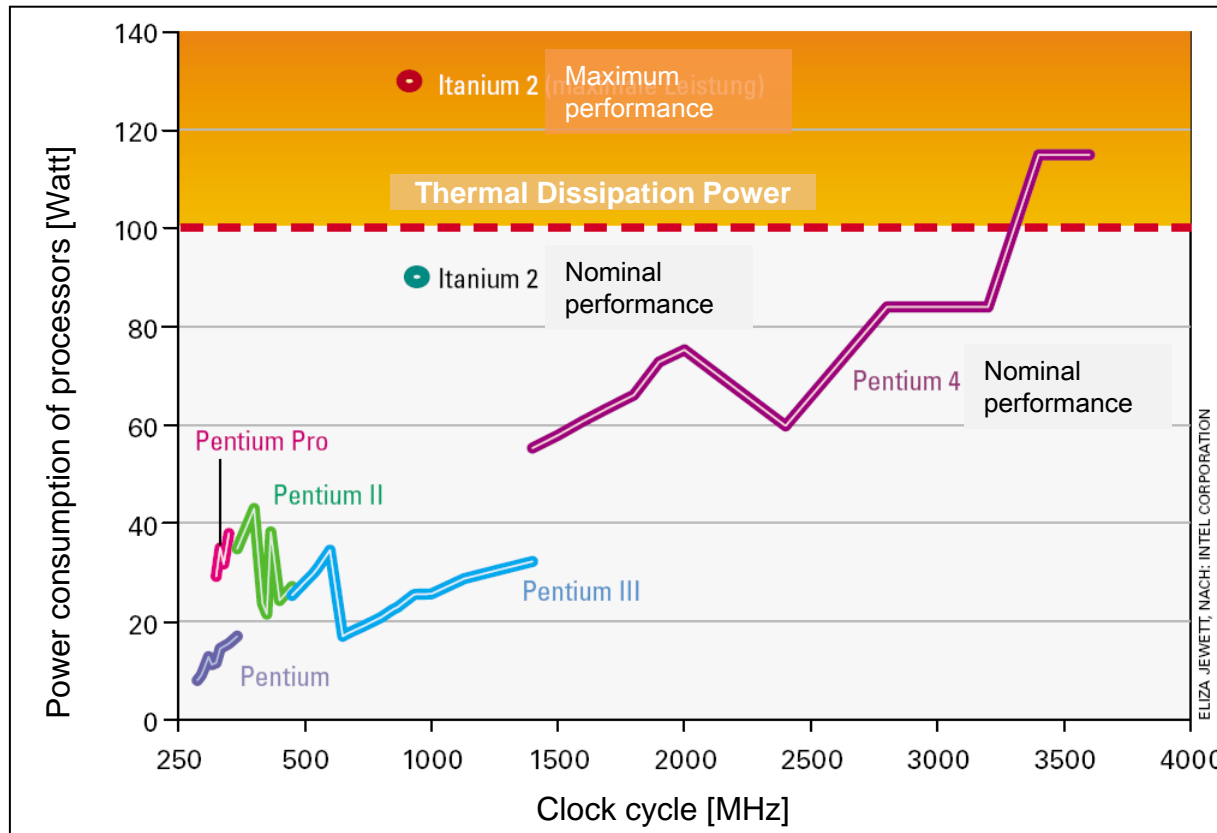
Clock wall

- Performance improvement until 2004
 - Primarily be increasing processor frequency
 - Building „better“ architectures was secondary
 - Continuous move to architecture side



The „Evolution“ to heterogeneous computing? Power wall

- Power Wall: Increasing the processor frequency increases the amount of required power



The „Evolution“ to heterogeneous computing?

Solution: Multi-core architectures

- Solution: Multi-core architectures
 - More processor cores on a single chip / processor package
 - Reason:
 - **Technology**: Increasing the clock frequency costs too much energy

$$P_{diss} = a \cdot f \cdot V_{dd}^2 \cdot C$$

Correlating the frequency f
and the supply voltage V_{dd}

Wayt Gibbs, Scientific American, 03/2005

„The question never was if, but when any why: When would processor manufacturers be forced to take it down a notch? Why would the iron law (cf. Moore’s Law) to produce a faster chip every two years no longer be sustainable?”

The „Evolution“ to heterogeneous computing?

Solution: Multi-core architectures

- Solution: Multi-core architectures
 - Reason:
 - **Architecture**: Principle of superscalarity almost exhausted
Dynamic Branch Prediction with more than 95% correctness

Justin R. Rattner, former CTO Intel

»We have the law of quadratic increase against us. We need exponential increase of the number of transistors with the side effects of increasing current consumption to achieve only marginal improvements for parallel instruction execution. «

Performance improvement could only be achieved using **„real parallelism“**

The „Evolution“ to heterogeneous computing?

Besides energy also further advantages

- Technological advantages of Multi-Core
 - Lower frequencies for individual cores
 - Even distribution of heat
 - Individual cores can be turned on/off or clocked according to needs
- Economic Advantages of Multi-Core
 - More than one core per die saves production cost
 - Additional processor core increases cooling requirements only linearly
- Architectural Advantages of Multi-Core
 - Previous architectural improvements almost exhausted

The „Evolution“ to heterogeneous computing? The road was not at the end

- Smaller, **specialized** cores → less overhead → less energy

• CISC / μ -programmed Processors




• RISC

• Superscalar Processors

• VLIW, EPIC, Multi-threading

• Multi-core / Many-core Processors

• Heterogeneous Architectures



Heterogeneous computing in embedded systems

Actually embedded was outrider

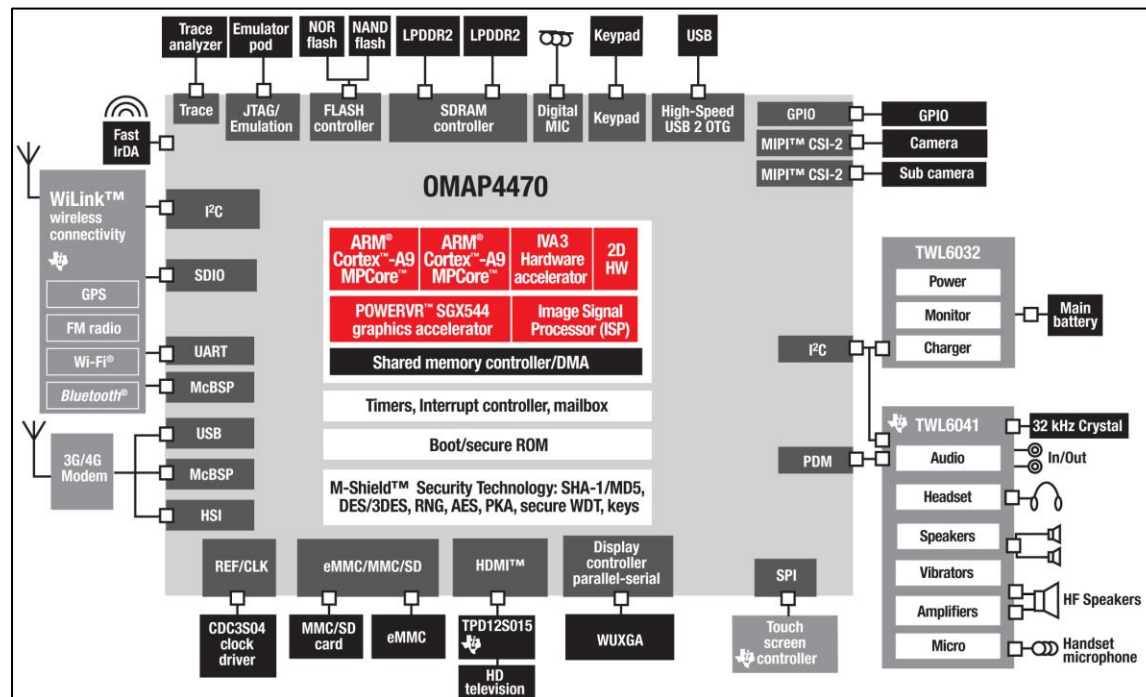
- Mobile systems have to save power
- Special core
 - Simpler pipeline structure
 - Simpler instruction set
 - Simpler control unit
 - More (specialized) compute power for less spent energy
 - Less flexibility → General purpose core still important

Heterogeneous computing in embedded systems

Example: Mobile phone processors



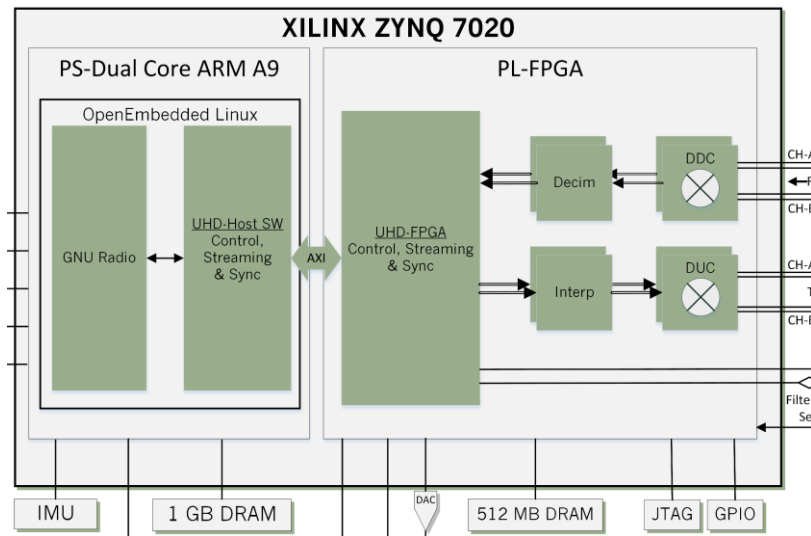
- TI OMAP
 - ARM Dual core processor as host
 - More specialized cores as accelerators
 - DSP + PowerVR GPU



Heterogeneous computing in embedded systems

Example: FPGAs



- Xilinx Zync
 - ARM Dual core processor as host + “glue logic”



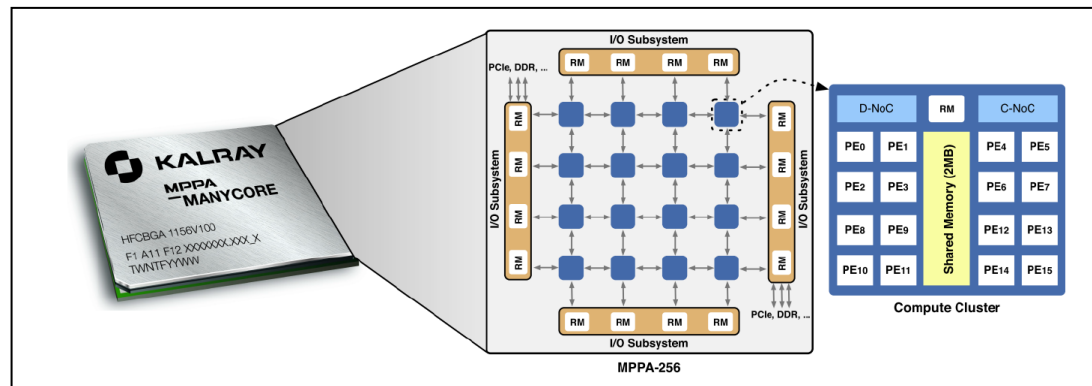
Heterogeneous computing in embedded systems

Example: Special many-core platforms

- Special many-core platforms
 - FPGA – SoC + Array of RISC processors from Adapteva



- SoC system MPPA from Kalray



Examples and challenges for embedded HIS

Heterogeneous architectures for smart cameras

Sensor fusion in automotive


Image Processing in Smart Cameras

Advantages

- Energy and space aware
- Low latency processing
- Embedded hardware sufficient

Architectures


- Dedicated architectures
- Real-time processing
- Heterogeneous architectures



Examples and challenges for embedded HIS

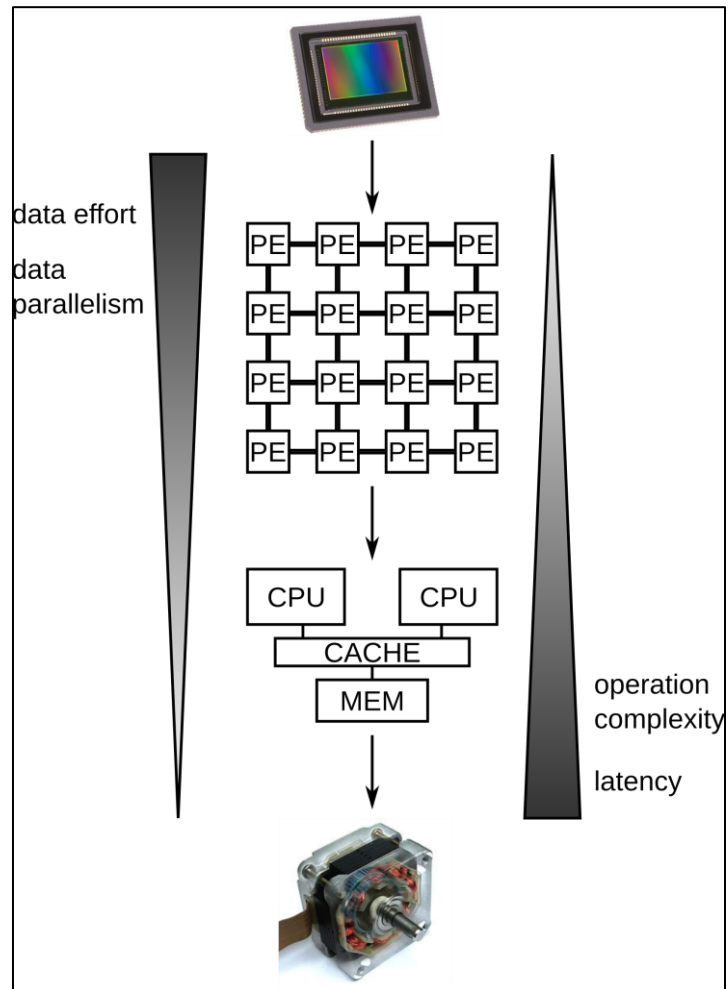
Heterogeneous Architectures for Smart Cameras

Higher energy cost for moving data than in-situ processing





Examples and challenges for embedded HIS

Heterogeneous Architectures for Smart Cameras



Examples and challenges for embedded HIS

Marker Detection



Examples and challenges for embedded HIS

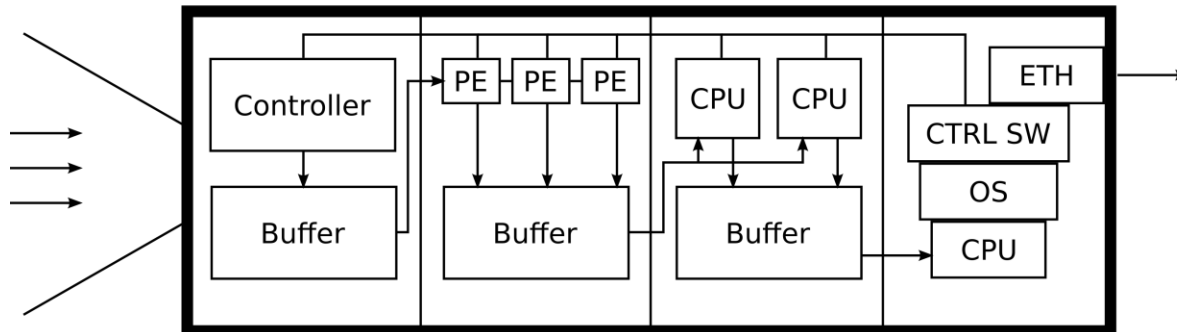
Heterogeneous Architectures for Smart Cameras

Sensorcontrol and
data acquisition



Image preprocessing
fine grained PEs

Image postprocessing
using CPUs

Systemintegration and
connection to the environment

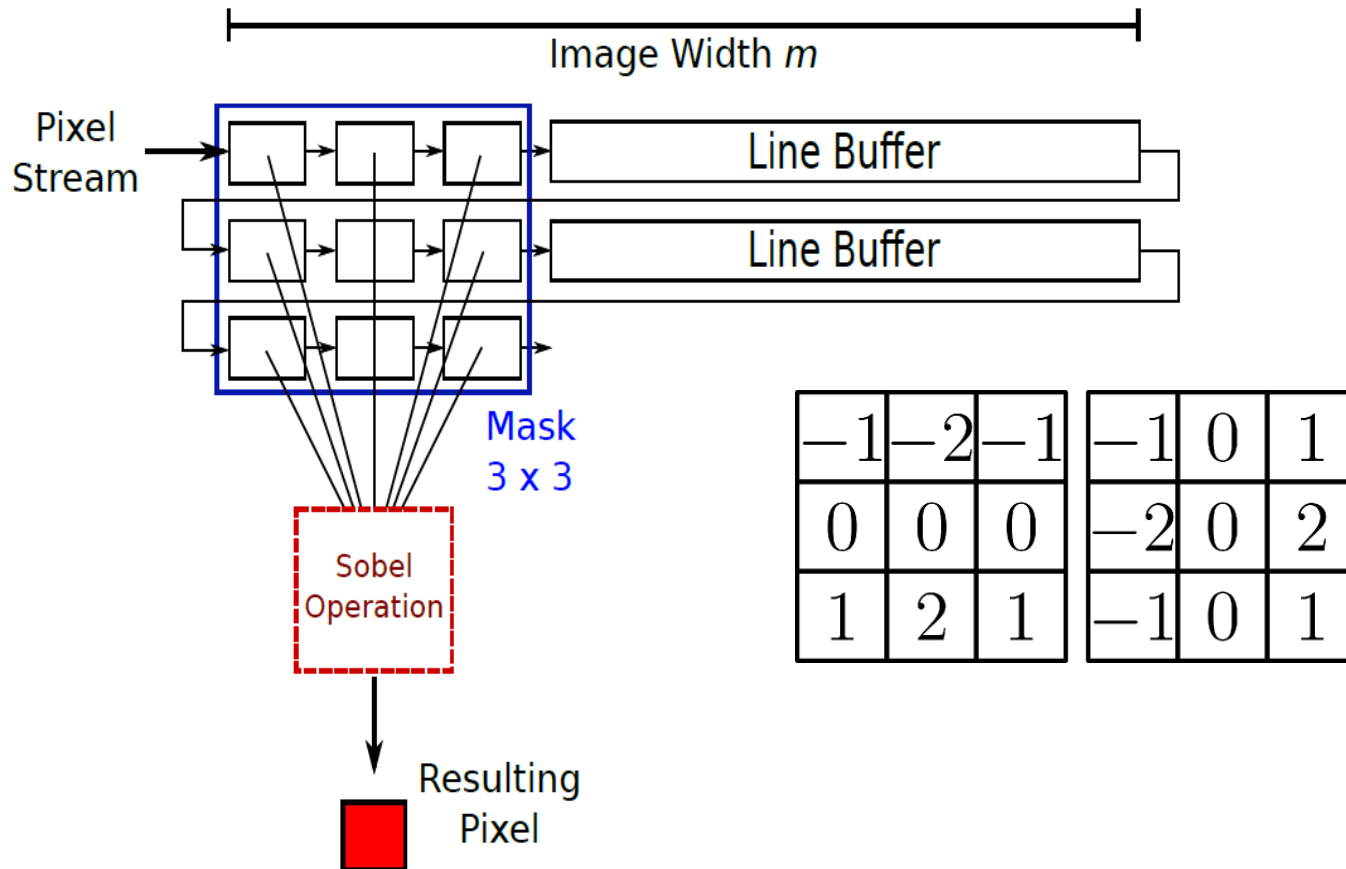



The Sensor




- CMOS VGA
 - 640 x 480 pixel
 - SCCB (I²C) interface
- Pre-processing on chip
 - White balance
 - Color correction
- 6 Euros on Amazon 😊


Edge Detection





Realization




Edge Detection



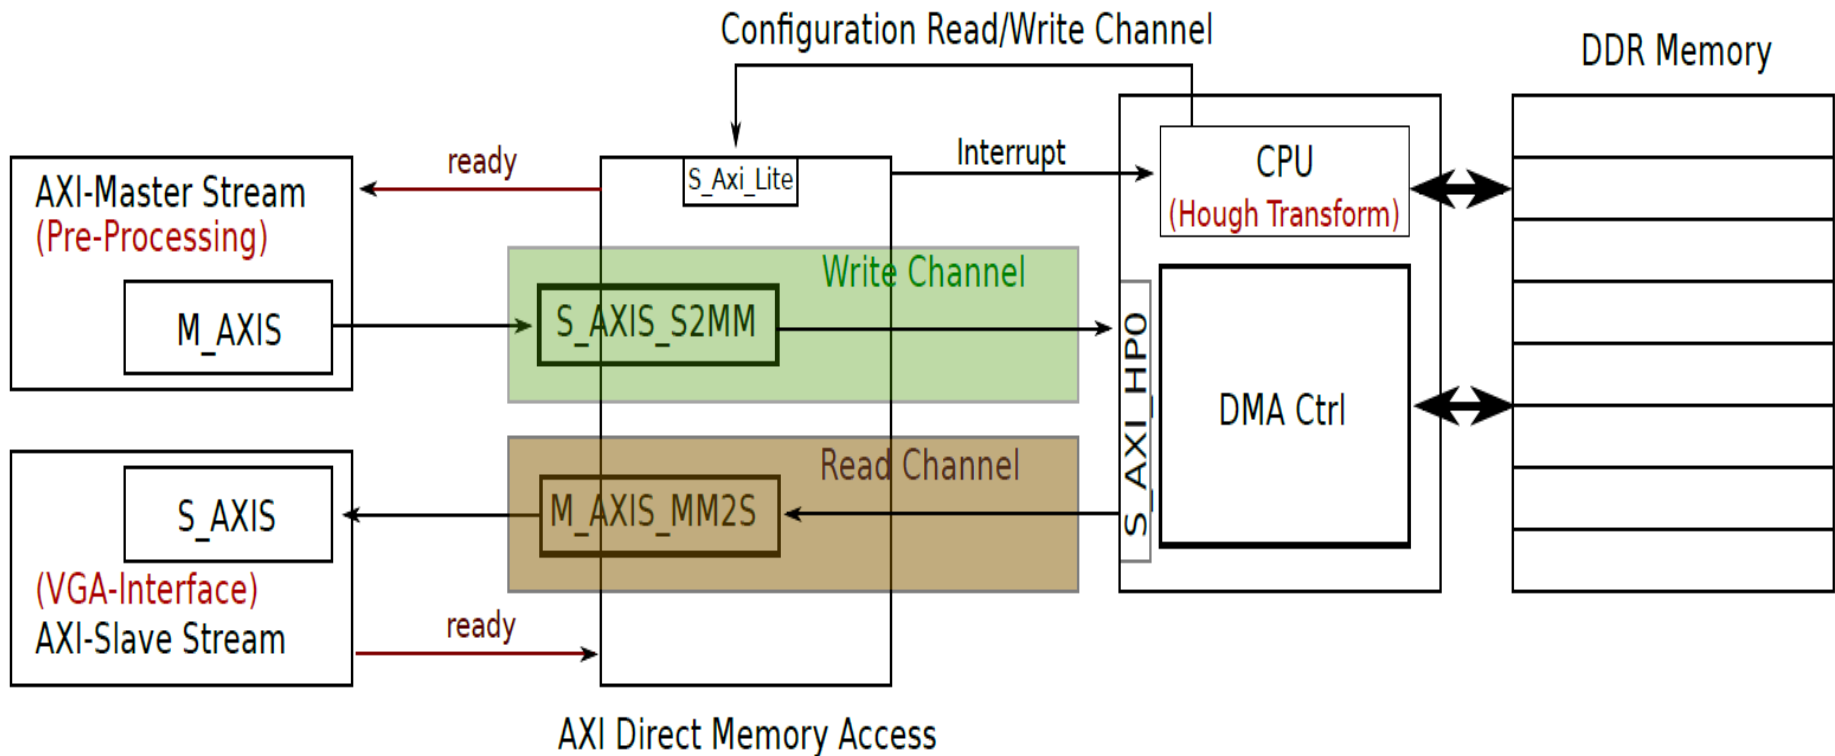
Circle Detection



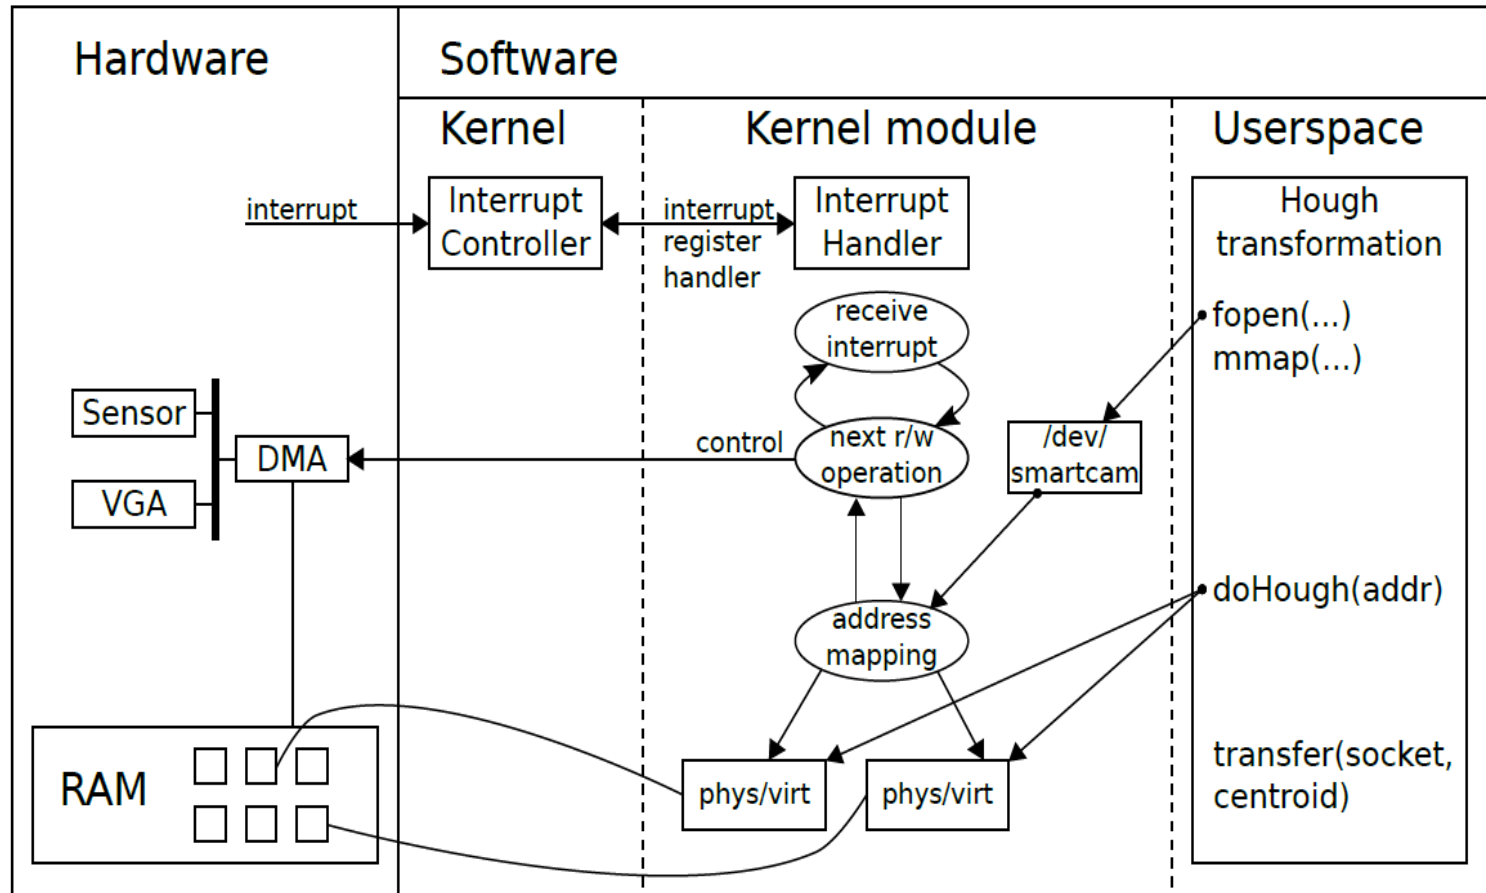

Circle Detection



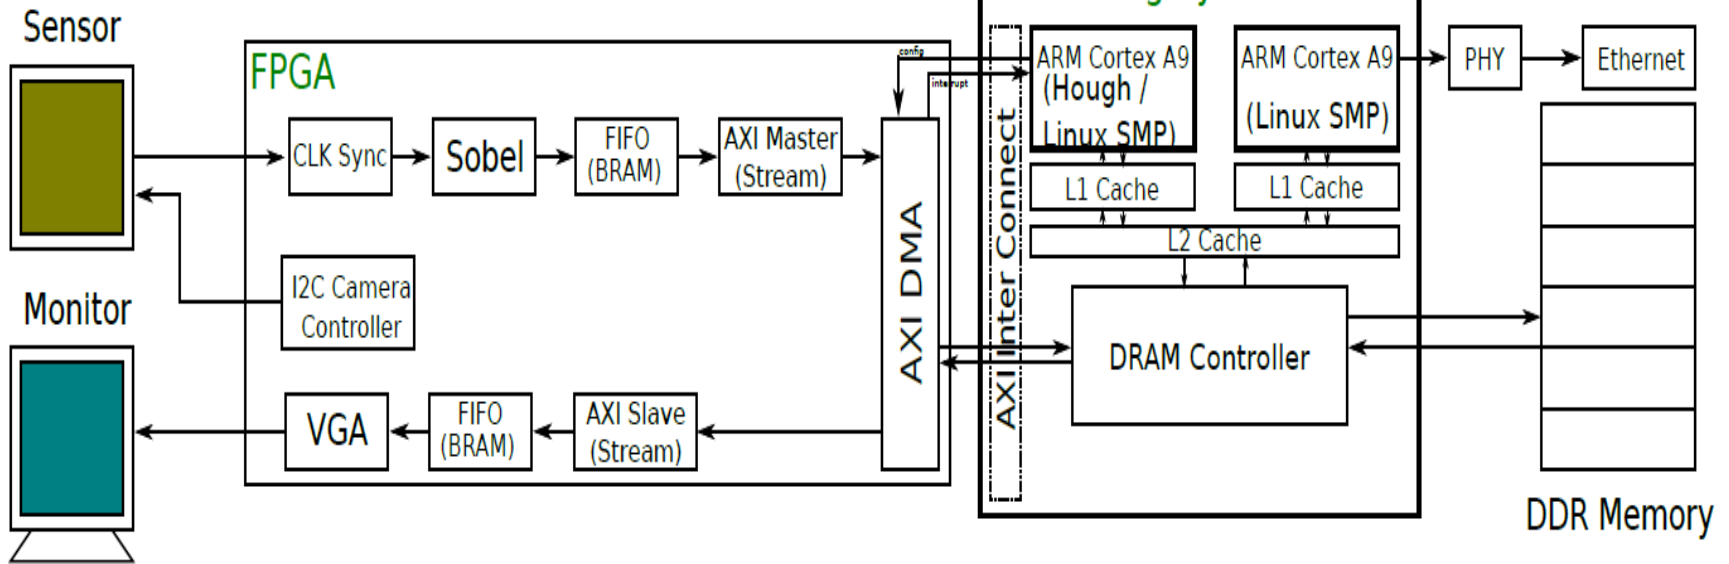

Using ARM Processors




System integration and environment



Complete Architecture



Working Prototype



Example: Robot Calibration



Examples and challenges for embedded HIS

Sensor fusion in automotive ADAS – Autonomous driving



Stanford University

Examples and challenges for embedded HIS

Time schedule to the way for Autonomous Driving

2016: Assisted driving

- Control of system still necessary

2020: Highly automated or piloted driving

- Driver has to take over control in difficult situations

2025: Complete autonomous driving

- Car is driving complete autonomously

Examples and challenges for embedded HIS

Tasks within an autonomous car

- Assistant to keep security distance / Cruise control
 - Front radar detects obstacles and other cars in front
 - Car adapts velocity
- Assistant to keep lane
 - Camera has to observe road marking
 - Alternative: infra red camera

Examples and challenges for embedded HIS

Tasks within an autonomous car

- Park distance control
 - Cameras and other sensors look out for surrounding obstacles and cars
- Navigation
 - Turning lanes, traffic lights and signs, road marking
 - Match GPS with road
- Construction of current environment
 - Prediction, estimation, security guarantee

Examples and challenges for embedded HIS

Sensors

- Radar
 - Distance and velocity measuring
- Lidar
 - Laser Scanner based on infrared
 - Laser Scanning: row sampling and image construction of the environment
- Ultra sonic sensors
 - Short range sensor to detect obstacles located aside the car
- Camera
 - Simple camera for detection of lanes and obstacles
 - Stereo camera for construction of 3D image for distance measuring and spatial vision (to replace front radar)

Examples and challenges for embedded HIS

Sensors

- Piloted driving today
- Rolling compute center in the boot
- Example: ZFAS components




Source: <https://www.audi-mediacycenter.com/de/pilotiertes-fahren-3651>



FORMUS³IC

Multi-Core safe and software-intensive Systems Improvement Community



AIRBUS

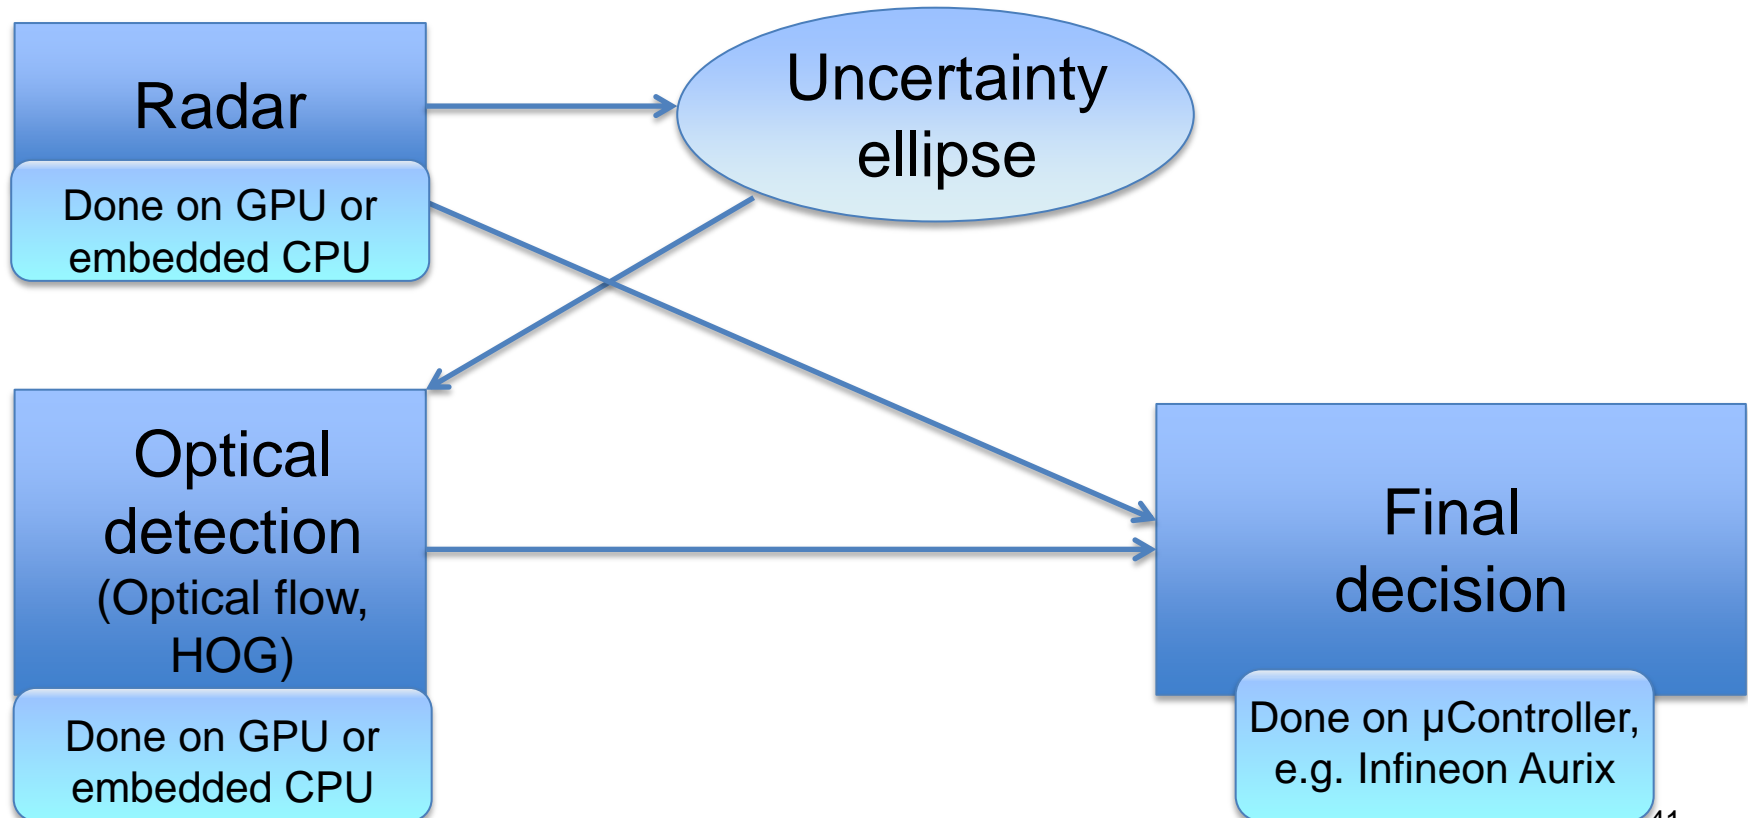


Audi



Examples and challenges for embedded HIS

Sensor fusion in automotive for ADAS – Autonomous driving



Examples and challenges for embedded HIS

Optical flow – Car detection & Pedestrian detection

- Moving of brightnesss at pixel (x,y)

$$I(x, y, t) = I(x + \delta x, y + \delta y, t + \delta t)$$

- Introduction of smoothness term introduced by Horn & Schunk

$$\begin{aligned}
 E(u, v) = & \iint (I_x + I_y + I_t)^2 \\
 & + \lambda \left[\left(\frac{\partial u}{\partial x} \right)^2 + \left(\frac{\partial u}{\partial y} \right)^2 + \left(\frac{\partial v}{\partial x} \right)^2 + \left(\frac{\partial v}{\partial y} \right)^2 \right] dx dy
 \end{aligned}$$

- Ending up in Euler-Langrage solution

$$I_x^2 u + I_x I_y v - \lambda \nabla^2 u + I_x I_t = 0$$


$$I_x I_y u + I_y^2 v - \lambda \nabla^2 v + I_y I_t = 0$$



Good for
GPU



Examples and challenges for embedded HIS

Optical flow – Car detection & Pedestrian detection



Examples and challenges for embedded HIS

HOG – Histogram of gradients algorithm



Examples and challenges for embedded HIS

Application Demo for HOG





HOG_DEMO(1).mp4



Examples and challenges for embedded HIS


Radar detection with beam-forming method



Examples and challenges for embedded HIS

Performance estimation for different architectures

SP FFT performance on Jetson TX1 (Quad-core ARM Cortex-A57, NVIDIA Maxwell GPU w. 256 CUDA Cores)



Challenges on software and design side for Heterogeneous Computing

Programming heterogeneous systems is a nightmare

The Yin and Yang of Heterogenous Hardware: Can Software Survive?
Prof. Kathryn McKinley, Microsoft Research, USA

- OpenCL → Performance loss
- Desirable → to predict performance
 - to predict what to calculate where
 - Performance prediction

Challenges on Software and Design Side for Heterogeneous Computing

Slides with TriCore / Aurix
topic are from Jens
Harnisch, Infineon, Munich

Infineon TriCore 1.6E and TriCore 1.6P

TriCore 1.6 Efficiency



Scalar Harvard

TriCore 1.6 Performance



Superscalar Harvard

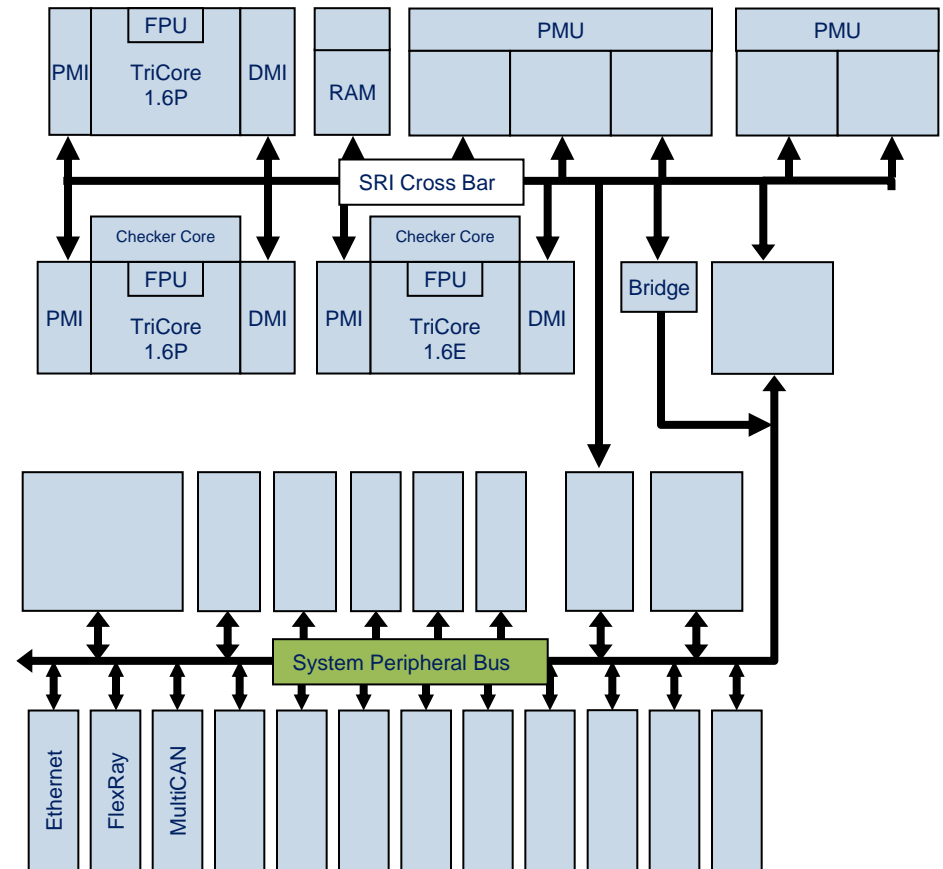
Common TriCore 1.6 Instruction Set
Single Precision Floating Point Unit - 2 FLOPs per cycle

- High **efficiency** architecture
- **4 pipeline stages** for up to 200MHz
- Power 0.2mW/MHz
- **1.2 – 1.4 DMIPS/MHz**

- High **performance** architecture
- **6 pipeline stages** for up to 300MHz
- High-performance DSP operations
- Power 0.3mW/MHz
- **1.6 – 2.3 DMIPS/MHz**

Challenges on software and design side for Heterogeneous Computing

- Highly predictable architecture with duplicated resources (local memories, crossbar) to avoid resource conflicts
- Starvation protection in crossbar
- No cache coherence: system predictability
- Dedicated and scalable communication instructions



Example Device

Support for hard real time systems

- Support for high average case performance usually contradicts high predictability
- Static timing analysis: precision of results and efficiency of analysis strongly depend on hardware architecture features
- Three classes of hardware architectures [1]:
 - Fully timing compositional architectures: no timing anomalies
 - Compositional architectures with constant-bounded effects: timing anomalies without domino effects -> TriCore is assumed to belong to this class
 - Non-compositional architectures

[1] Predictability Considerations in the Design of Multi-Core Embedded Systems.

C. Cullmann, C. Ferdinand, G. Gebhard, D. Grund, C. Maiza, J. Reineke, B. Triquet, S. Wegener, and R. Wilhelm. *Ingénieurs de l'Automobile*, 807, 2010.



Design Guidelines for predictable multicore architectures



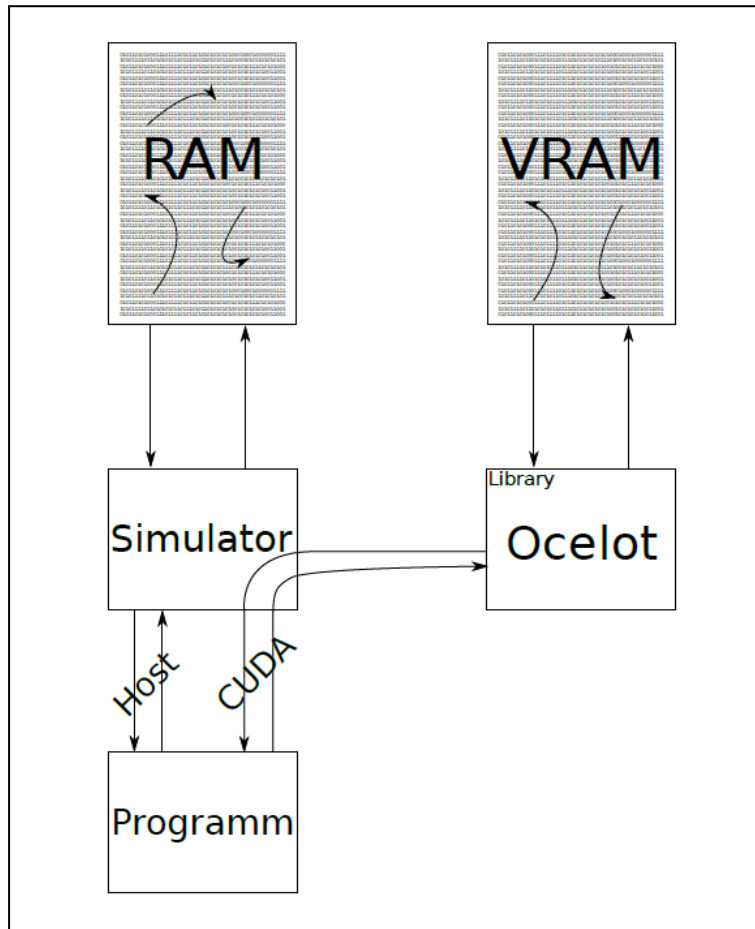
- Established by C. Cullmann, C. Ferdinand, G. Gebhard, D. Grund, C. Maiza, J. Reineke, B. Triquet, S. Wegener and R. Wilhelm [1]

1. Fully timing compositional architecture ✓
2. Disjoint instruction and data caches ✓
3. Caches with LRU replacement policy ✓
4. A shared bus protocol with bounded access delay ✓
5. Private caches ✓
6. Private memories, or, only share lonely resources (depends on software) ✓

[1] Predictability Considerations in the Design of Multi-Core Embedded Systems. C. Cullmann, C. Ferdinand, G. Gebhard, D. Grund, C. Maiza, J. Reineke, B. Triquet, S. Wegener, and R. Wilhelm. Ingénieurs de l'Automobile, 807, 2010.

Challenges on software and design side for Heterogeneous Computing

What about GPU?



Source: Nvidia.com

Conclusion

- Route to Heterogeneous Computing was driven by Energy
- In particular important for Embedded Systems
- Smart Camera:
 - Driven by energy and space constraints: pre-processing data better than moving
- Sensor fusion in Automotive
 - Driven by performance and energy
- Critics:
 - HIS are not good due to circumstantial Software Programming (see Cell Processor)
- Thesis
 - Even if hard: Better design and programming environments are needed
 - Due to energy constraints: Wheels will not be turned back

**Thank you very
much!**



**FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG**

TECHNISCHE FAKULTÄT