Karlsruhe Institute of Technology

FPGAs in Detector Instrumentation: concepts and trends

Michele Caselle and Marc Weber

12-15 March 2019 Physics Department, TU Dresden

Helmholtz Research Field Matter

Challenges for the future

The next generation of detectors are extremely challenging: HEP, astrophysics, photon science, etc.

Simulation of 10 000 tracks in a future HL-CMS Tracker detector

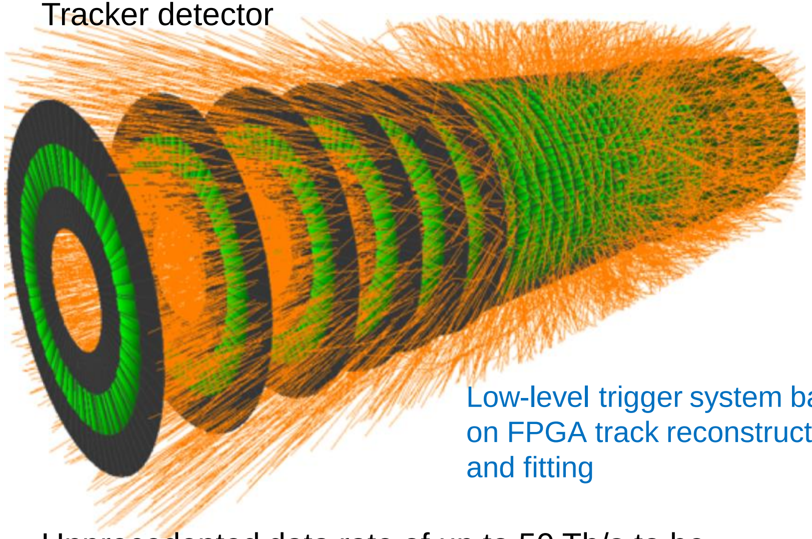

Low-level trigger system based on FPGA track reconstruction and fitting

Unprecedented data rate of up to 50 Tb/s to be processed in < 4 µs with high efficiency and acceptable fake rate

High brilliance coherent synchrotron source for new generation of photon science experiments: EuXFEL, FLASH, FLUTE, KARA, SLS-2, etc.

Advanced beam diagnostic tools for new synchrotron and plasma accelerator machines.

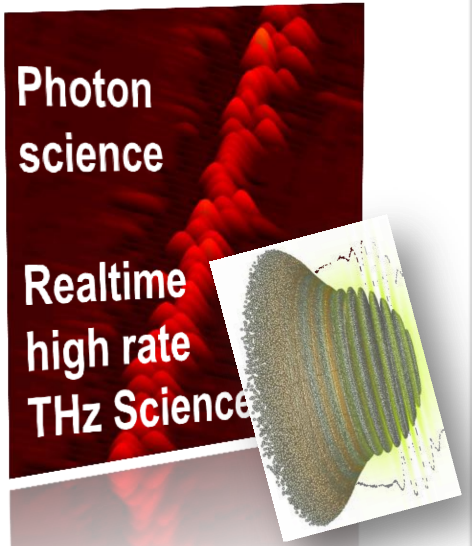

Complex dynamics on short time scale, terahertz detector technologies, sub-ps time resolution, O(100) Gb/s continuous data acquisition for long-time scale (sec- hrs)

Scaling of existing technologies is necessary to be prepared for future experiments

Challenges for the future (II)

- Increased luminosity requires:
  - Higher segmentation → O (10 µm) on hybrid and down to O(1-3 µm) on monolithic
  - Higher hit-rate capability → sub-nanosecond
  - Higher radiation hardness

Tracker will be replaced for better precision and high rates & radiation → large pile-up (HL-LHC) → **4D tracking** (timing within pixel layers)

DAQ system, interface to, and control, front-end (‘*readout*’) & organize data into coherent structure (‘*event building*’) to cope with enormous data volume

Trigger systems are essential to find rare events & new physics (‘*trigger*’)

Configuration and control of detector (‘*run control*’)

Provide monitoring of system and data quality (‘*DQM*’)

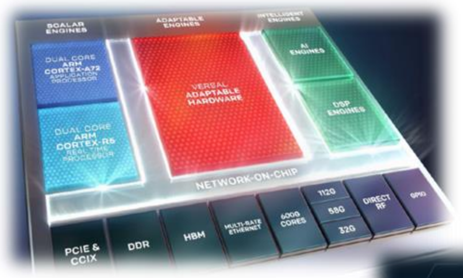

Next generation of FPGAs **(?)**

Next generation of data processing **(?)**

Next generation of FPGA

A New Class of Devices for Today's Challenges

MPSoC (Zynq US+, 2016)
1.5 GHz Quad-core ARM processor **+**
FPGA @ 890 MHz F_{MAX} (all register used)

MPSoC (Zynq US, 2014)
1 GHz Dual-core ARM processor **+**
FPGA @ 740 MHz F_{MAX} (all register used)

VERSAL architecture

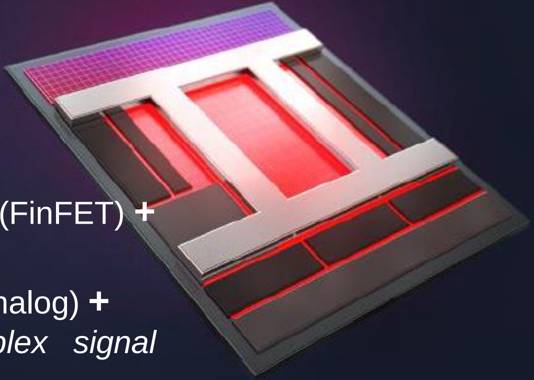

FPGA based on 7 nm TSMC (FinFET) **+**
Multi-core processors **+**
Radio frequency front-end (analog) **+**
Artificial intelligent & complex signal processing (hard-core)

SoC (Virtex-4, 2005)
FPGA @ 300 MHz F_{MAX} **+**
450 MHz PowerPC processor

Invented 1985
(XC2064)

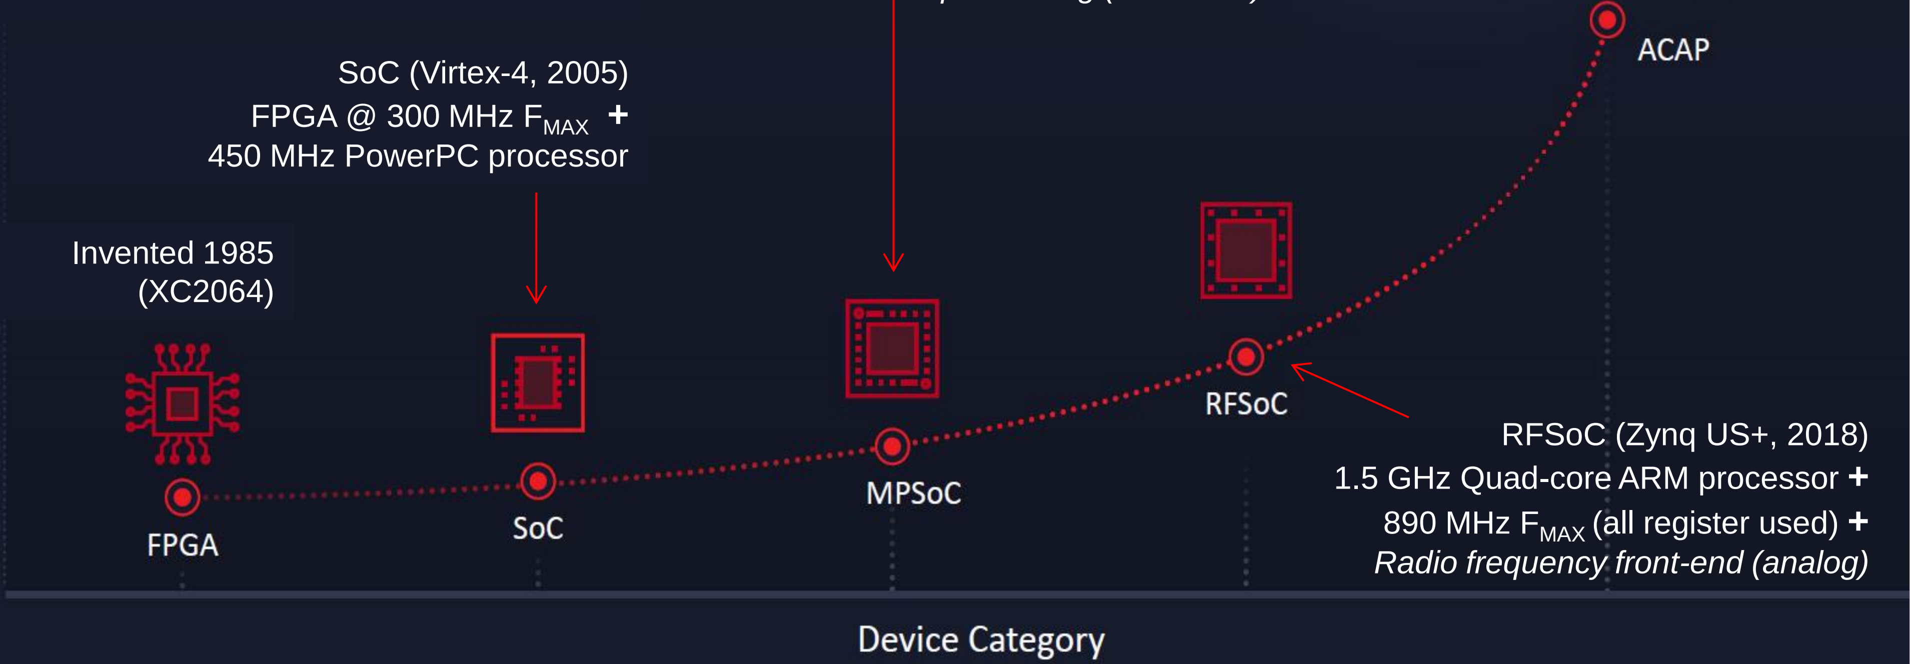

RFSoC (Zynq US+, 2018)
1.5 GHz Quad-core ARM processor **+**
890 MHz F_{MAX} (all register used) **+**
Radio frequency front-end (analog)

Three Ages of FPGA

Age of Invention

- FPGAs are much smaller than the application problem size
- No tools, manual P&R

Age of accumulation

- Communication is main market
- Dedicated logic blocks (e.g. high-speed I/O, multiplier)
- Low cost FPGAs
- IP cores for large FPGAs

Age of heterogeneous

1984 1992 2000 2008 2014

Age of expansion

- Rapidly growing FPGA sizes
- Increasing Demand for Design Automation
- SRAM FPGAs first for new technology – domination

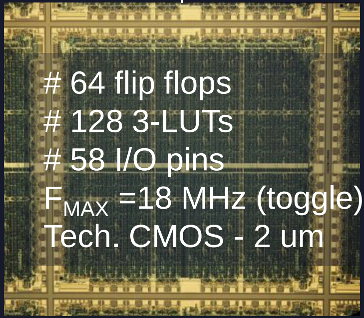

XC2064 first FPGA (1985)

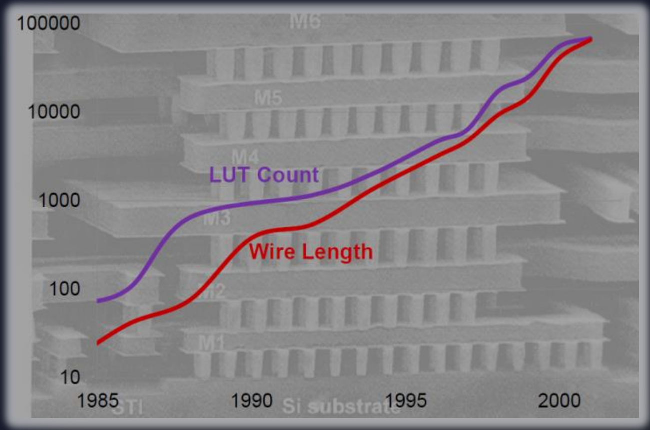

Heterogeneous: ZYNQ MPSoC technology

Heterogeneous platform of the Zynq System-on-Chip (SoC) integrates, in a monolithic device, FPGA resources with a back-end software running on a hard-core ARM-based processor.

Quad-core ARM processor

Dual-core real-time processor

High-flex2 DAQ card based on ZYNQ US+. Developed @ KIT

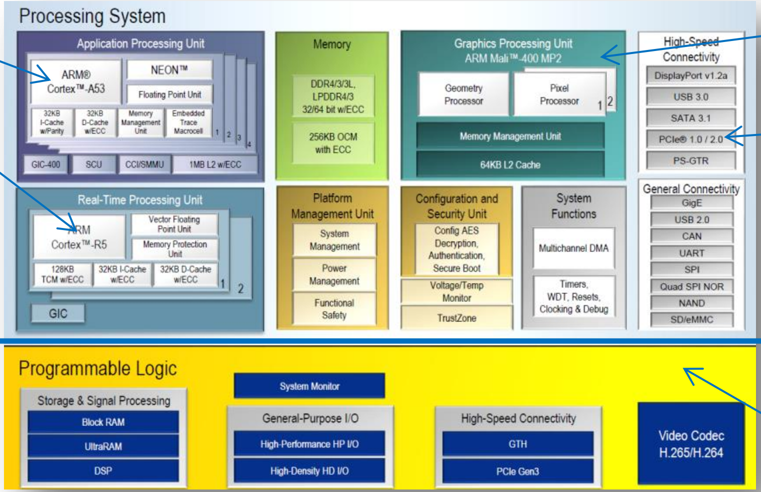

Embedded Graphics Processing Unit (GPUs)

Standard peripherals hard IP-core

Multi-Processor System-on-Chip (PS)

FPGA programmable logic (PL)

FPGA & High-speed interfaces based on TSMC 16 FinFET

- Next talk: *System-on-Chip FPGAs: Experience and Recommendations* by Ralf SPIWOKS *(CERN)*
- System-on-Chip (SoC) workshop – CERN , 12-14 June 2019 https://indico.cern.ch/event/799275/registrations/48961/

Ref. https://www.xilinx.com/support/documentation/selection-guides/zynq-ultrascale-plus-product-selection-guide.pdf

Heterogeneous: ZYNQ *RF*SoC technology

The bandwidth bottleneck of previous ZYNQ MPSoC was the bandwidth and complexity of the JESD204B standard communication for fast ADC/DAC

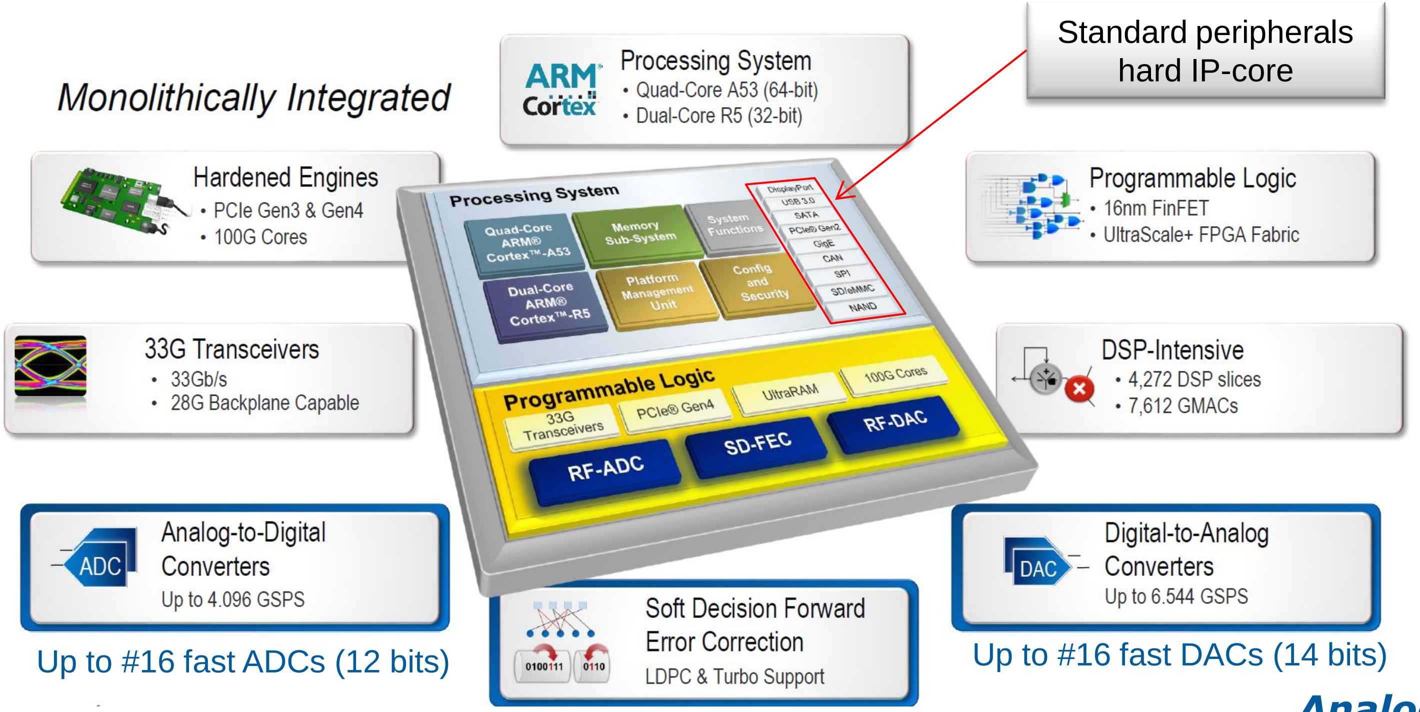

Up to #16 fast ADCs (12 bits)

Up to #16 fast DACs (14 bits)

Analog layers

Ref. https://www.xilinx.com/support/documentation/selection-guides/zynq-usp-rfsoc-product-selection-guide.pdf

RF-ADC/DAC Implementation steps

- Add RF-ADC/DAC instance using the IP-Core
- Use GUI to configure and customize the analog block
- Connect the ADC/DAC to ZYNQ by "Running_connection_automation"

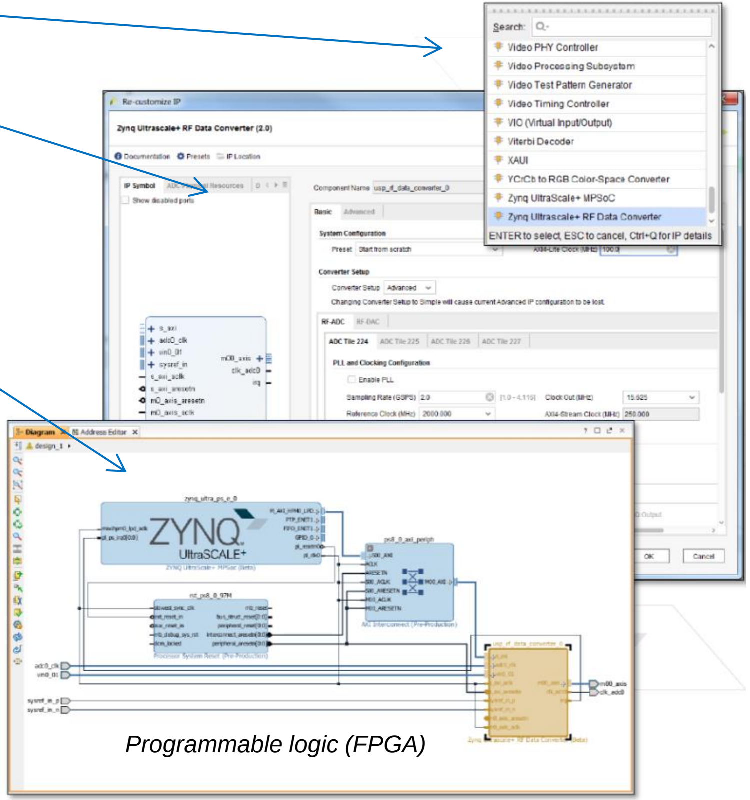

Direct RF Sampling & digital Signal Processing

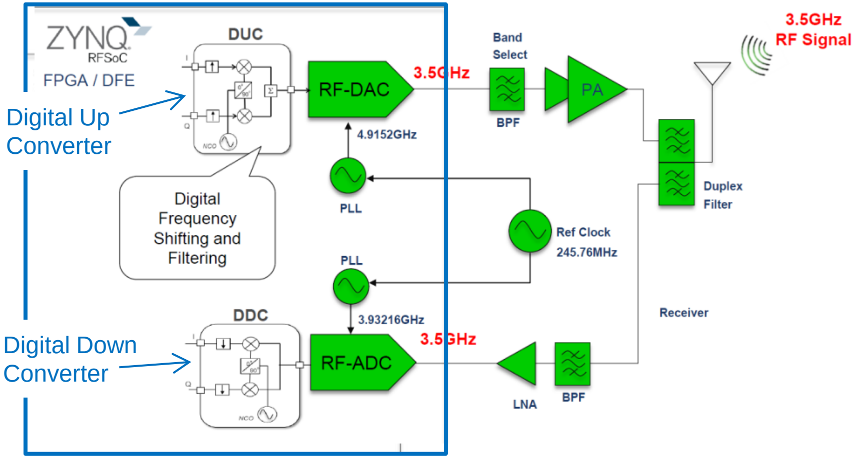

Versal architecture - Overview

New Xilinx architecture **ACAP** (Adaptive Compute Acceleration Platform) develop in TSMC 7nm FinFET technology, key features:

- More heterogeneous
- **Processor + FPGA + RF analog** + **HBM** + **AI engines** (artificial intelligence)
- High bandwidth & low-latency → multi-Terabit/sec throughput

Scalar processors + real time

Adaptable Engines (FPGA)

Vector processors AI + DSP Engines

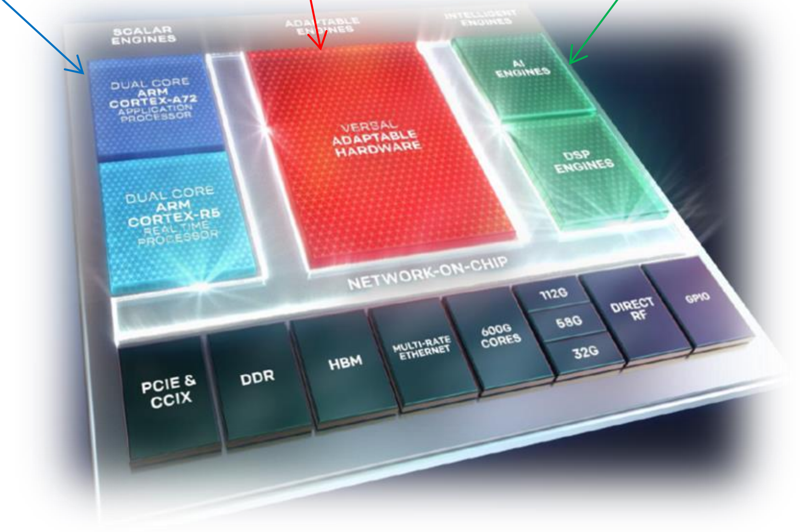

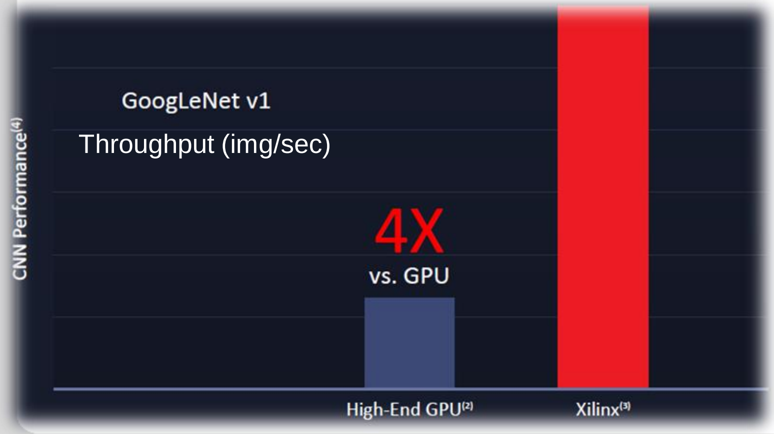

AI Inference acceleration

Leveraging AI engines

Ref: https://www.xilinx.com/support/documentation/white_papers/wp505-versal-acap.pdf

Next generation of data processing

Next generation of data processing

New generation of particle physics experiment already reached hundreds of millions of events per second, meaning physicists must sift through tens of *petabytes of data per year*. As the resolution of detectors improve, ever better solution is needed for real-time pre-processing and filtering (‘trigger’) of the most promising events.

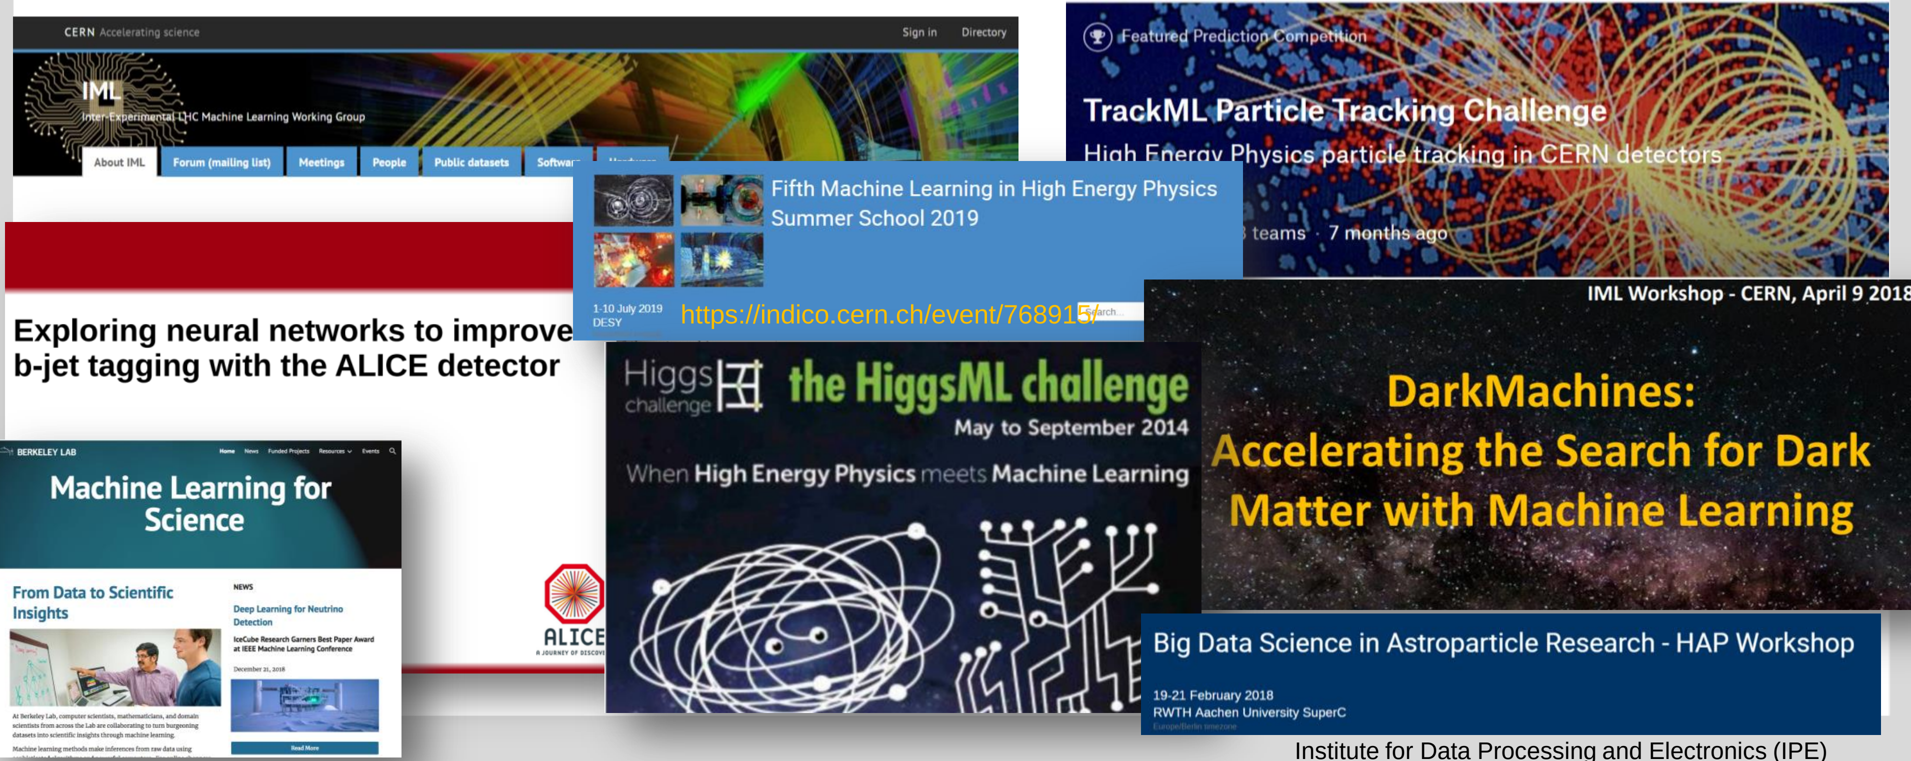

What’s the different between Artificial Iintelligence, Machine Learning, and Deep Learning?

AI involves machines that can perform tasks that are characteristic of human intelligence
While this is rather general, it includes things like planning, understanding language, recognizing objects and sounds, learning, and problem solving

At its core, machine learning is simply a way of achieving AI

“*the ability to learn without being explicitly programmed*.” You see, you can get AI without using machine learning, but this would require building millions of lines of codes with complex rules and decision-trees.

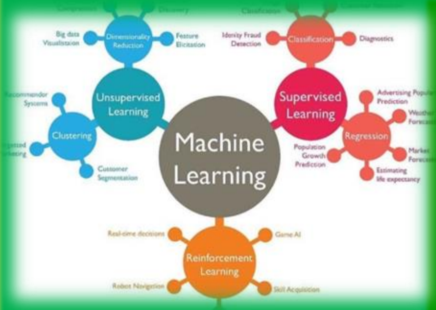

Deep learning is one of many approaches to machine learning

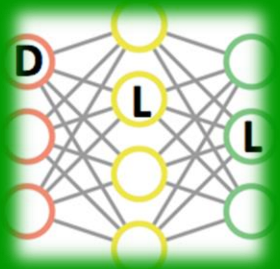

Deep learning was inspired by the structure and function of the brain, namely the interconnecting of many neurons. Artificial Neural Networks (ANNs) are algorithms that mimic the biological structure of the brain.

Machine learning

- Data analysis
- Reconstruction chain → jet tagging (task to find the particle ID of a jet)
- Low-level and high-level trigger systems
- Data quality monitor
- etc.

Machine learning

- Physics analysis
- Reconstruction chain → jet tagging (task to find the particle ID of a jet)
- Low-level and high-level trigger systems
- Data quality monitor
- etc.

Using heterogeneous FPGA-GPU to combine the strengths of both FPGA and GPU technologies

- Heterogeneous FPGA/GPU-based readout system

http://ufo.kit.edu/ufo

Data source

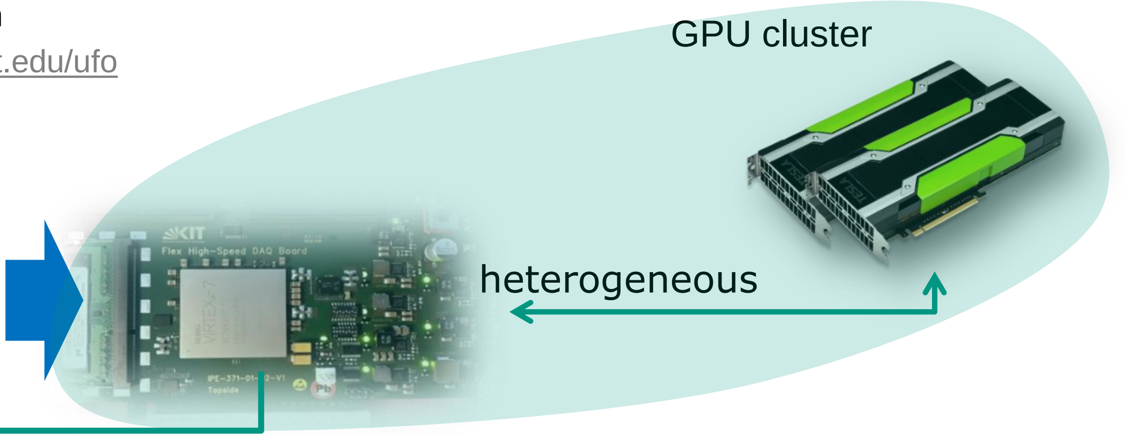

Fast feedback

M. Caselle et al., JINST 12 C03015 (2017)

Machine learning

- Physics analysis
- Reconstruction chain → jet tagging (task to find the particle ID of a jet)
- Low-level and high-level trigger systems
- Data quality monitor
- etc.

Using heterogeneous FPGA-GPU to combine the strengths of both FPGA and GPU technologies

Reducing the time needed for training

FPGA (ML *Inference*)
Low-latency, customized designs, high-performance memories resource, many interfaces, broader flexibility, etc.

GPU cluster

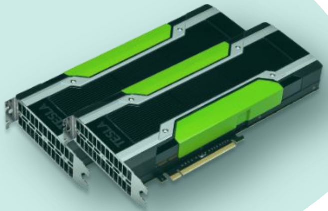

GPU (ML *training*)
floating point and large processing capabilities, easy to program, etc.

Data source

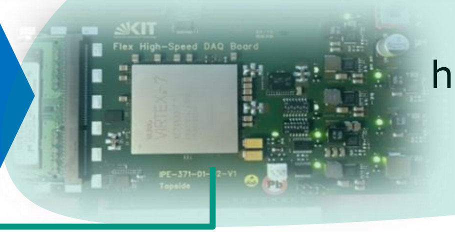

heterogeneous

Fast feedback

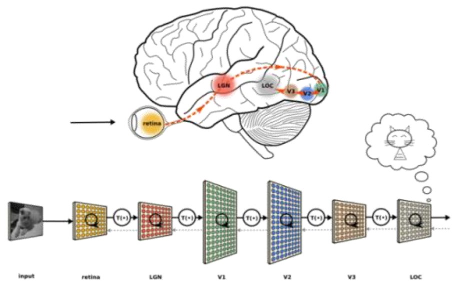

M. Caselle et al., JINST 12 C03015 (2017)

Machine learning

- Physics analysis
- Reconstruction chain → jet tagging (task to find the particle ID of a jet)
- Low-level and high-level trigger systems
- Data quality monitor
- etc.

Using heterogeneous FPGA-GPU to combine the strengths of both FPGA and GPU technologies

FPGA (ML *Inference*)
Low-latency, customized designs, high-performance memories resource, many interfaces, broader flexibility, etc.

GPU cluster

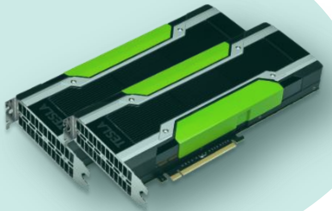

GPU (ML *training*)
floating point and large processing capabilities, easy to program, etc.

theano
Caffe2
TensorFlow
Keras

Data source

heterogeneous

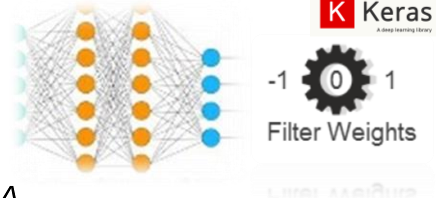

Fast feedback

DNN with weight are pruned, synthetized and implemented on FPGA

Machine learning

- Physics analysis
- Reconstruction chain → jet tagging (task to find the particle ID of a jet)
- Low-level and high-level trigger systems
- Data quality monitor
- etc.

Using heterogeneous FPGA-GPU to combine the strengths of both FPGA and GPU technologies

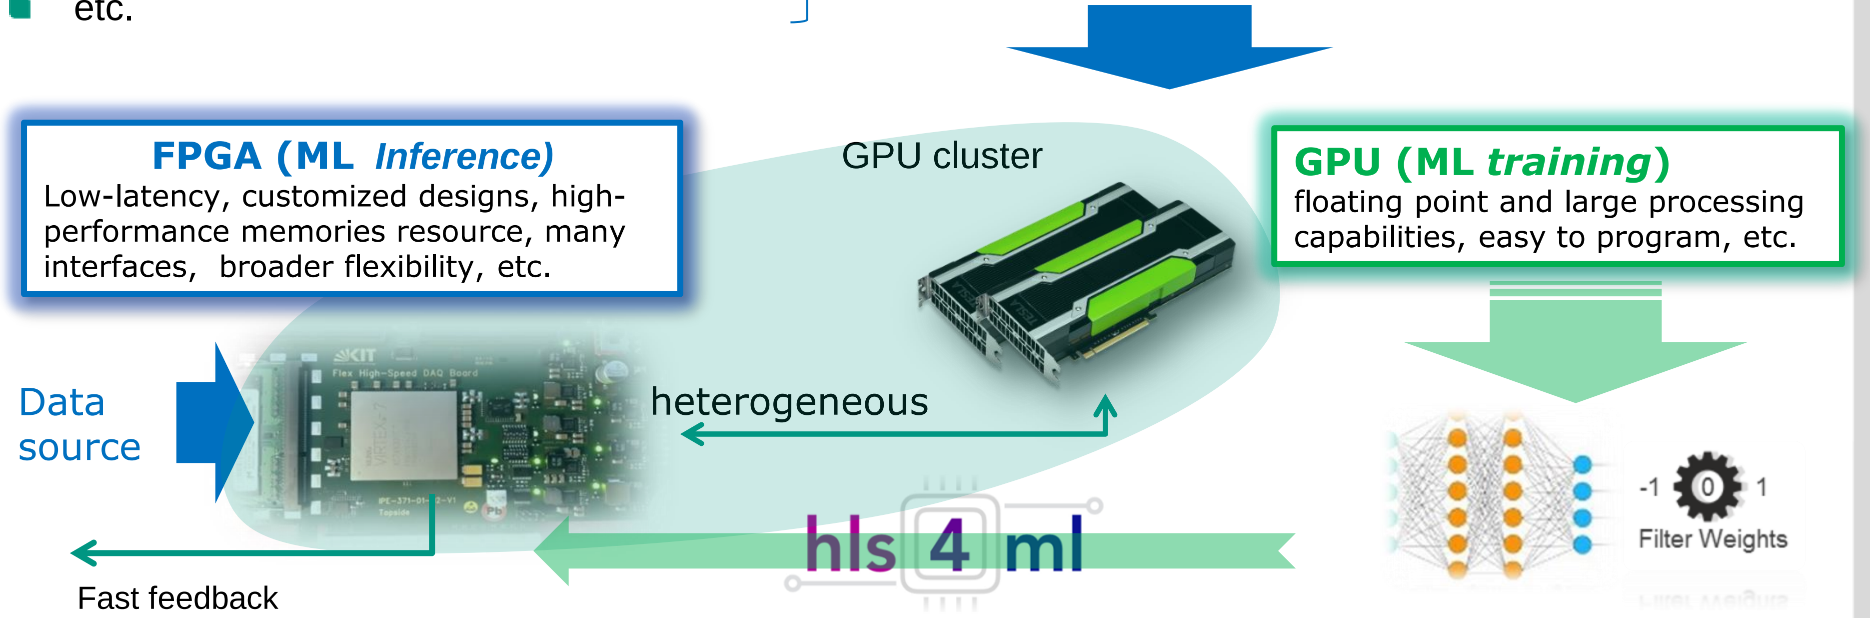

High-performance Deep Machine Learning FPGA

High Level Synthesis four Machine Learning has been developed at CERN

Talk: Machine Learning for HL-LHC Detectors, by Jennifer NGADIUBA (CERN), Friday room GER 038

What is ***hls4ml*** → framework removes major barrier on hardware development of ML algorithms allowing developers with little or no FPGA expertise to program FPGA by high level synthesis tools

Case studies:

- **HEP**: Machine Learning for low-level trigger system based on FPGA → fast jet substructure classification
- **Photon Science**: fast feedback to RF-system of synchrotron machine for the control of the coherent THz emission in burst mode
- *Under development* the low-level trigger system for HL-CMS track finding

Jet classification

Res_V	FPGA prediction	Python Keras calculation (GPU)
Gluon (g)	0.118164	0.12993355
Quark (q)	0.639648	0.6487177
Boson (W)	0.118164	0.10633943
Boson (Z)	0.118164	0.10616959
top quark (t)	0.015625	0.00883975

16 clock cycles @ 200 MHz → latency = 80 ns
Implemented on ZYNQ UltraScale+ XCZU9EG

Low-latency & ultra-fast ML inference

Javier Duarte et al., Fast inference of deep neural networks in FPGAs for particle physics, arXiv:1804.06913v3 (2018)

Deep Machine Learning – photon science

Karlsruher Institut für Technologie

Goal is to keep the coherent THz intensity stable by fast feedback to RF System

Complex and nonlinear dynamics in longitudinal / transverse bunch profiles → described by a nonlocal nonlinear partial differential ***vlasov fokker planck*** equation describing the time evolution of the probability distribution of a particle in synchrotron machine

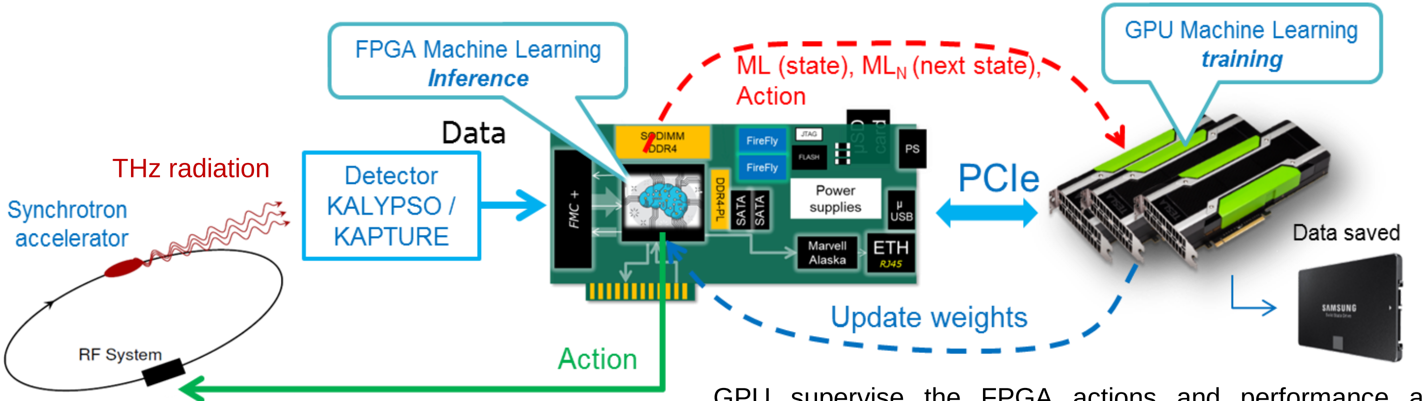

GPU supervise the FPGA actions and performance a continuous training → updating the FPGA weights

5 layers of fully connected layers neural network

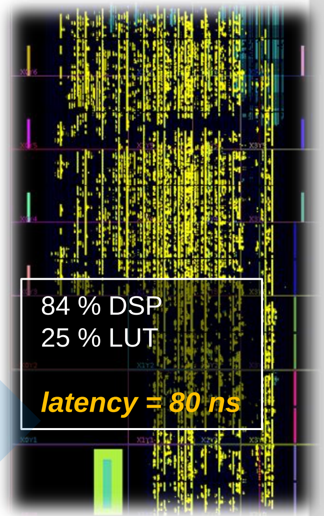

Keras
A deep learning library
Training on GPU
hls4ml
VIVADO
HLx Editions

Impact

Karlsruher Institut für Technologie

- Next generation of detectors will be harder than HL-LHC, not necessary ‘bigger’, but will generate huge and complex data volume (petabyte/sec), serious technological challenges → we are required to push the technology envelope (especially for trigger and data acquisition)

- Toward a common hardware infrastructure → novel programmable devices families, i.e. ACAP, which combine **high-performance FPGA**, **scalar processors** + **artificial intelligence** + **vector processors** (i.e. GPU)

  - Ad-hoc logics → in the high performance & flexible FPGA region
  - Detector/system interconnections → by flexible I/O and multi-Tb/s integrated high bandwidth memory
  - AI engine → low-latency high-performance ML inference (trigger, processing)

- Machine learning is one of the most promising technique for data analysis, processing (trigger), intelligent detector configuration, data quality monitor, etc.

  - Very **flexible,** the ***ability to learn*** without being explicitly programmed, simple **code maintenance, easy integration** in FPGA, etc.

… we are just at the beginning, new exciting technologies will appear in the next years

Thank you for your attention

Backup slides

FPGA market or which vendors did survive?

Karlsruher Institut für Technologie

achronix SEMICONDUCTOR CORPORATION FPGA IP core for SoC designs	intel FPGA SRAM based FPGAs Broad range	LATTICE SEMICONDUCTOR Low power & cost efficient FPGAs
MICROCHIP 2 small SRAM FPGA families	Microsemi Flash & Antifuse FPGAs	XILINX ALL PROGRAMMABLE SRAM based FPGAs Broad range

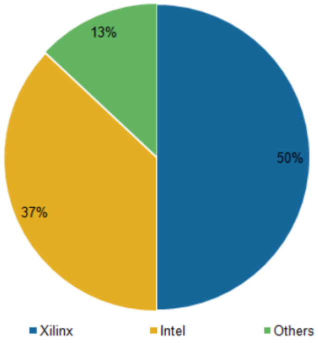

Xilinx: 50 %

Altera / Intel: 37 %

Others: 13 %

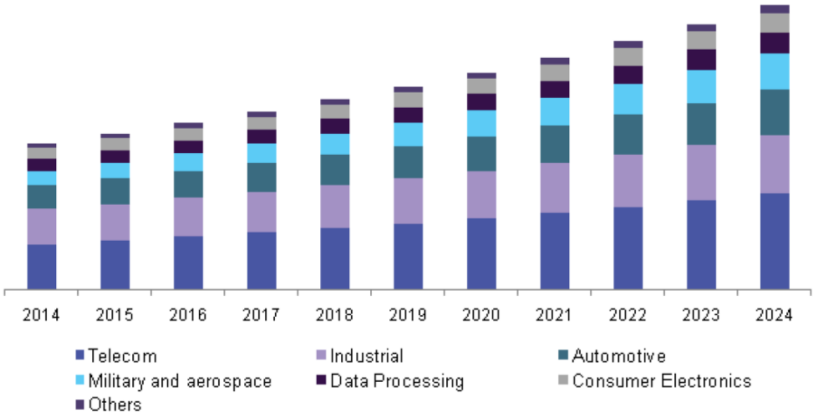

Two FPGA technology solutions

Karlsruher Institut für Technologie

Virtex 7T

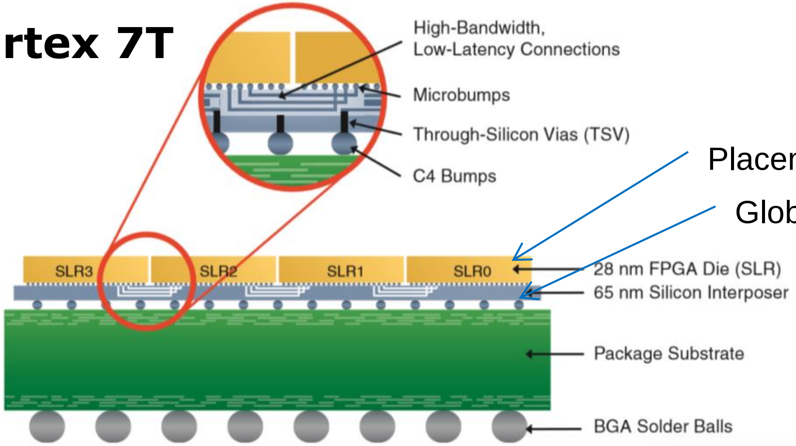

Placement & local routing

Global routing

Virtex 7T FPGA Enabled by SSI (Stacked Silicon Interconnect) technology

DDR bandwidth GAP because DDR on board

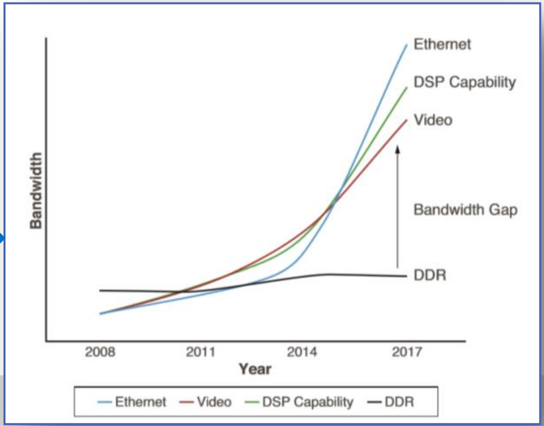

Solution

Virtex Ultrascale+

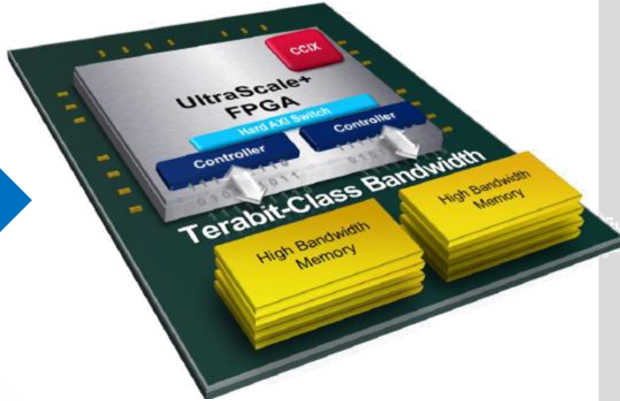

High Bandwidth Memory integrated into the FPGA

- DRAM: up to 4 GByte
- Bandwidth: 1.84 Tb/s

Ref. Xilinx white paper WP380

Charting an Aggressive Course Forward

Courtesy of Hanak, Georg (Avnet Silica)

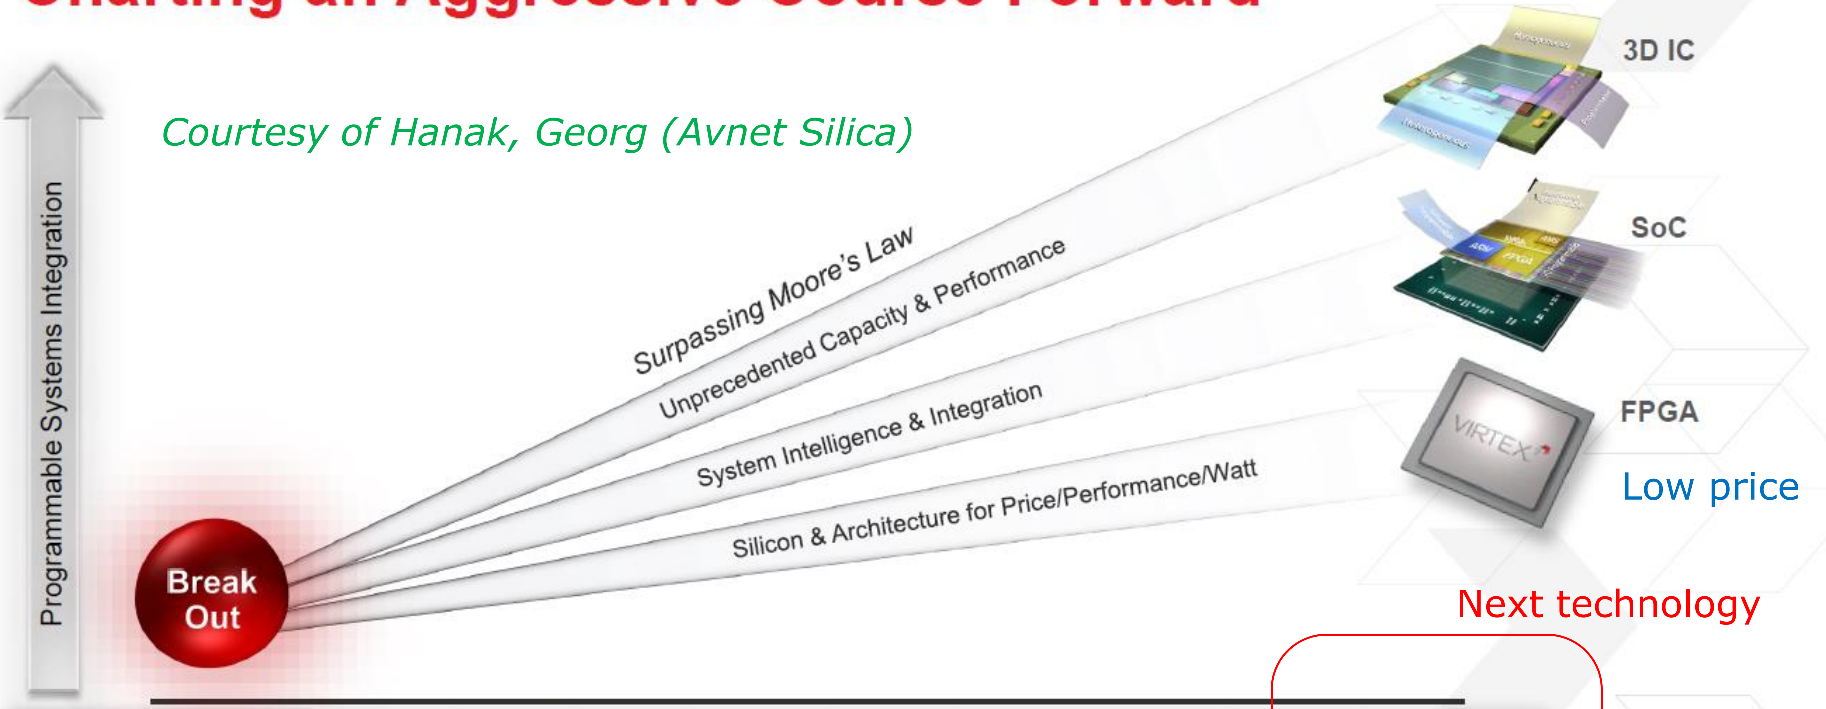

Next technology

Tools	VIVADO	UltraFAST™ Design Methodology	SDx	**VERSAL**
Architecture	7 Series	UltraScale™ Architecture		**ACAP**
Process	28HPL	20SoC	16FinFET+	7nm

XILINX.

High-performance Deep Machine Learning (II)

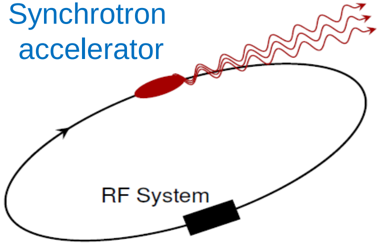

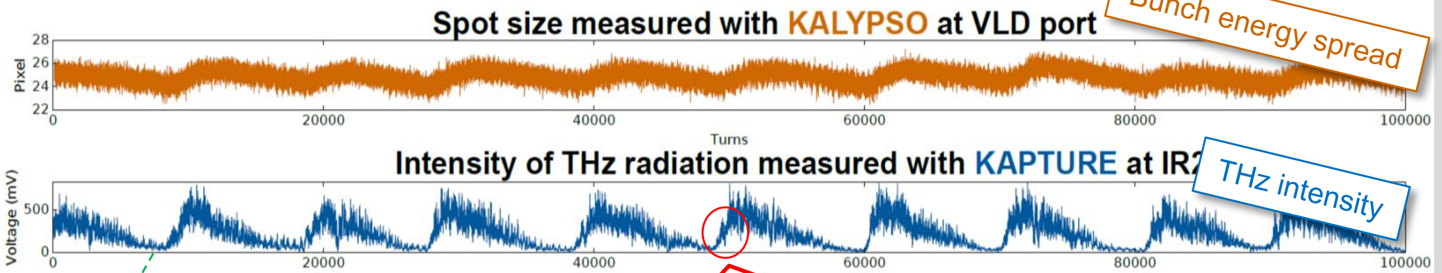

Longitudinal bunch profile measured with KALYPSO

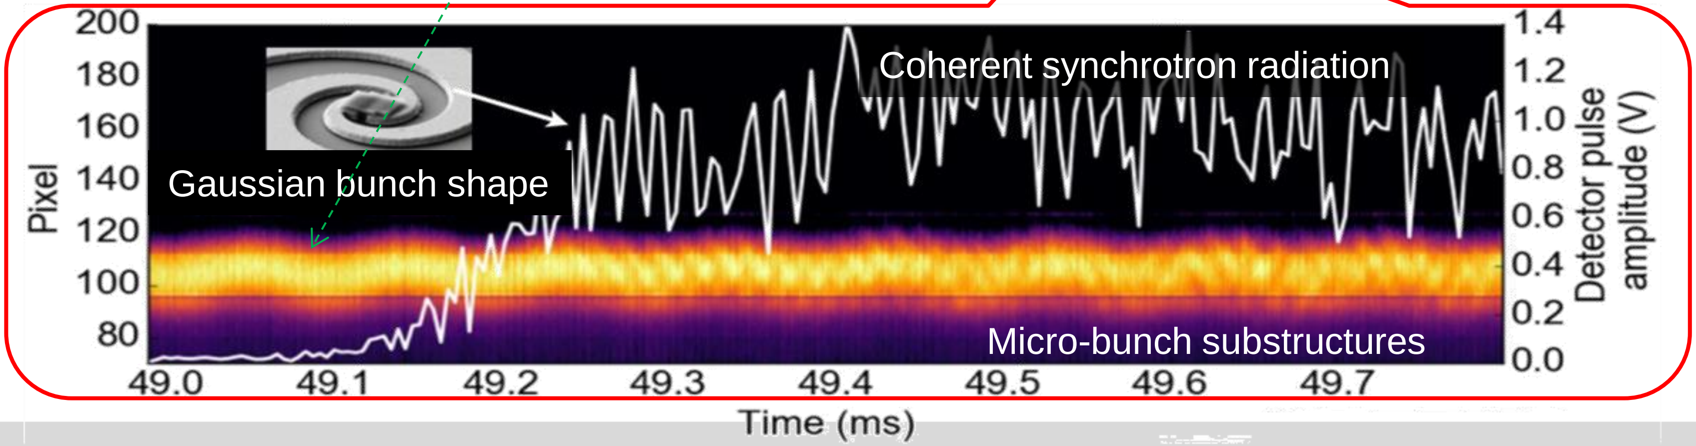

New “High-Flex 2” readout card

- **Photon science:**
  - New generation of ultra-fast cameras for science (PS ⇄ DTS)
  - Detectors for beam diagnostics (ARD ⇄ DTS)

    → Poster by M. Caselle (ARD + DTS)
  - Hardware platform for implementing machine learning algorithms
- **Superconducting sensors and quantum technologies:**
  - Readout of Metallic Magnetic Calorimeters arrays
  - Control- and readout of superconducting quantum bits
- **High Energy Physics (HEP):**
  - NA62 (SPS-CERN) fast “low-level” trigger system, GPU-based ~µs latency
  - High Level Trigger (HLT) based on modern GPUs and FPGAs accelerators

Optical link (full-duplex) ***up to 190 Gb/s***

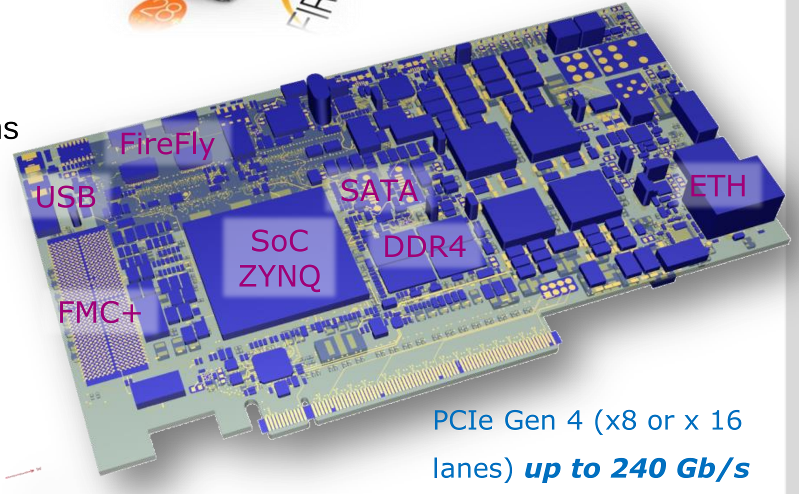

PCIe Gen 4 (x8 or x 16 lanes) ***up to 240 Gb/s***

(MU ⇄ DTS)

New "High-Flex 2" readout card

- **Photon science:**
  - New generation of detectors for beam diagnostics
  - Diagnostics and stabilization of laser systems
- **Superconducting sensors and quantum technologies:**
  - Control- and readout of superconducting quantum bits
- **High Energy Physics (HEP):**
  - NA62 (SPS-CERN) fast "low-level" trigger system, GPU-based
  - High Level Trigger (HLT) based on GPU- FPGAs accelerators
- **Hardware platform for Artificial Intelligence algorithms**
  - Heterogeneous FPGA- GPU system based on Machine learning

Fully designed @ KIT

New Readout card

Karlsruher Institut für Technologie

Optical link: 12 lanes @ 28 Gb/s (full-duplex) > 330 Gb/s.

FMC+ 57.4 connector:
32 transvers @ 28 Gbps
160 lines @ 2 Gbps

µUSB + UART connectors

SODIMM-DDR4 for PS

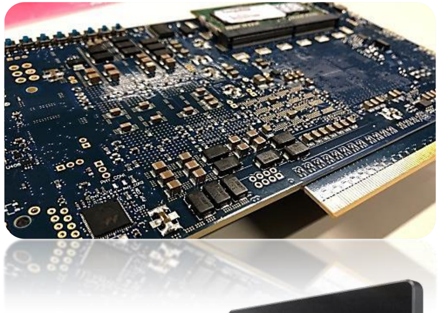

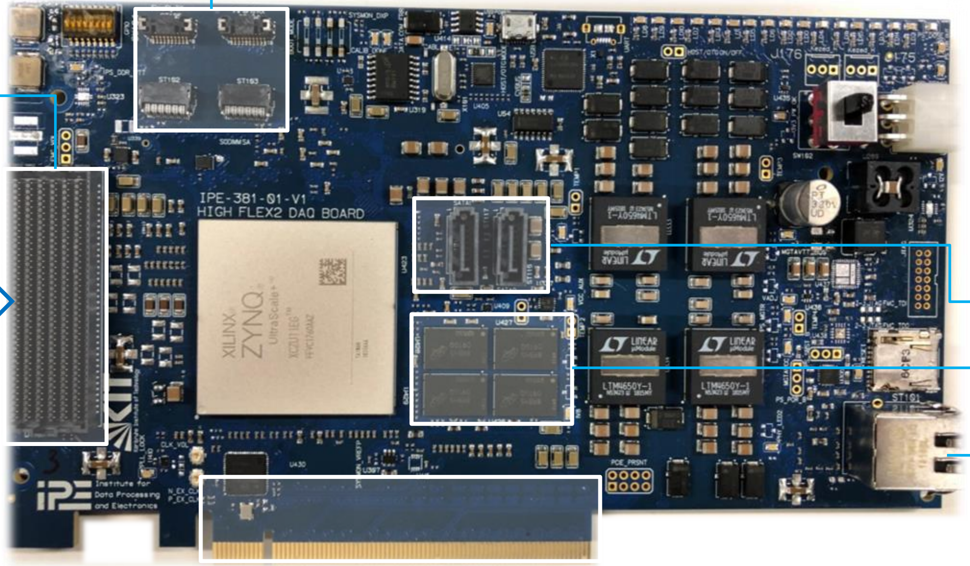

Detector

Two **SATA** connections:
Data saved directly on SSDs

DDR4 for PL (FPGA)

ETH connected to processor (TCP/IP)

PCIe Gen 3 and 4, 16 lanes → up to 240 Gb/s full-duplex

Modularity/scalability readout architecture

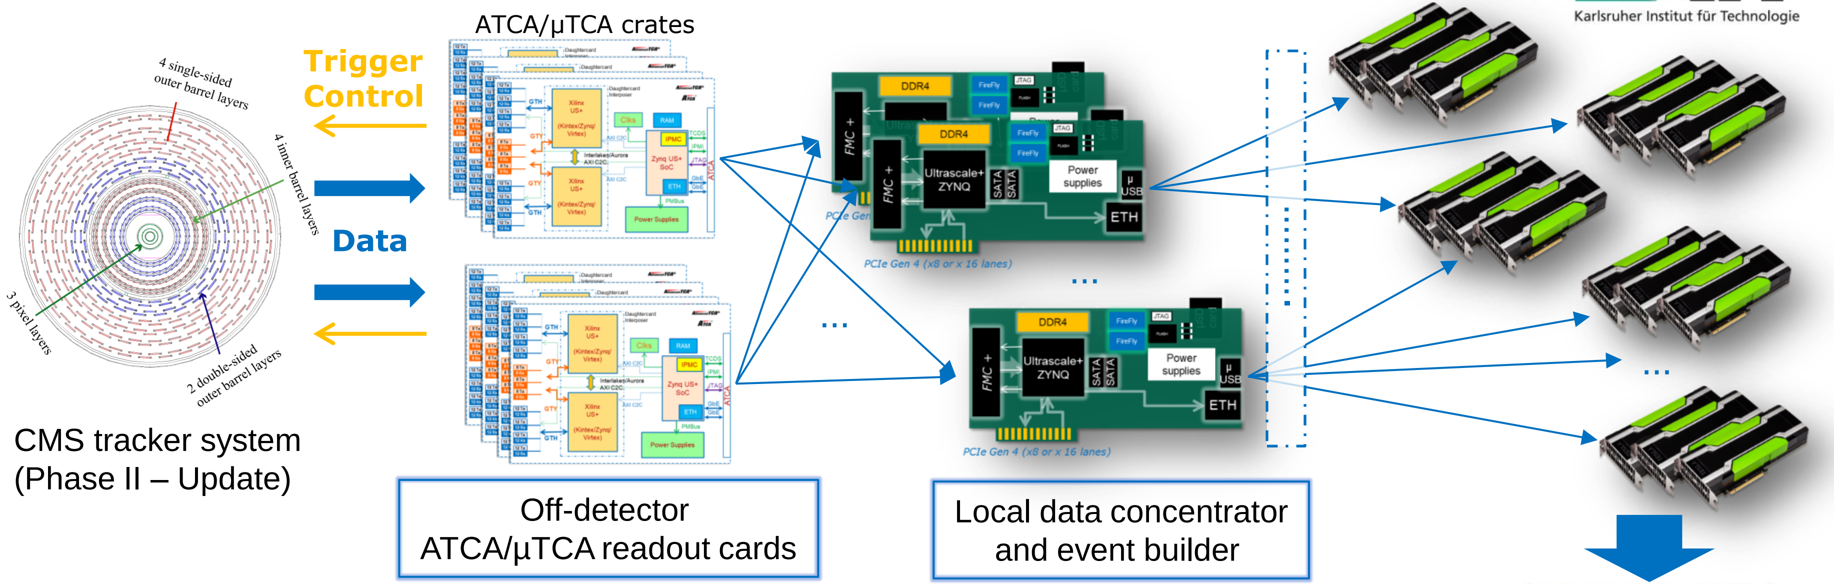

Key HW/SW developments:

- ATCA/µTCA readout card for HEP and photon science
- Heterogeneous FPGA-GPU architectures based on PCIe readout cards
- Machine learning algorithm for HEP (in collaboration with Cern)

High Level trigger
and event display

Readout card based on ATCA / µTCA

- ATCA or µTCA standard
- Compatible with CMS timing layer detector readout card
- Key ideas:

Finisar

1. I/O by interposer:

- Optical links
- FMC+ digital
- LPAF analog + digital

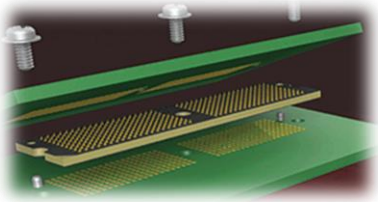

High-density board-to-board interconnection by interposer technology

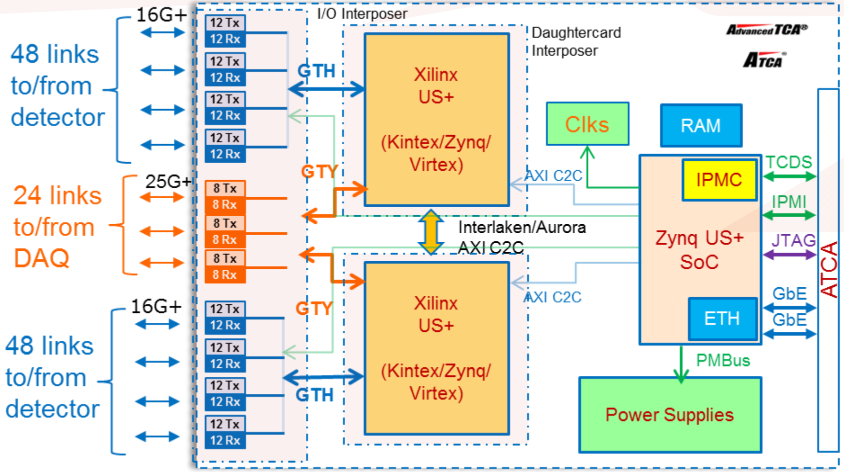

2. FPGA by interposer:

- High-flexibility in FPGA selection
- Kintex/Virtex or RF/SoC Zynq

3. Centralized slow-control by ZYNQ processors and ATCA shelf manager, new tasks:

- Detector calibration
- On-line data quality check
- Machine learning

Common readout card 3 in 1 concept

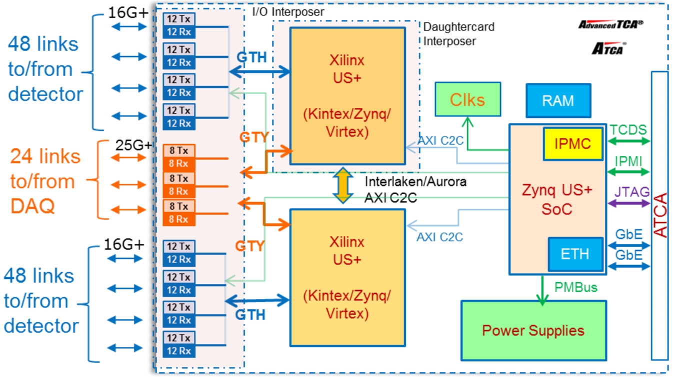

Common HW for:

Massive optical I/O → 96 + 24 optical links

High-end ADC/DAC cards → ADSCs = 80 GS/s + DACs = 320 GS/s

Massive digital connection → 320 digital I/O + 40 high-speed serial link (each > 16 Gb/s)

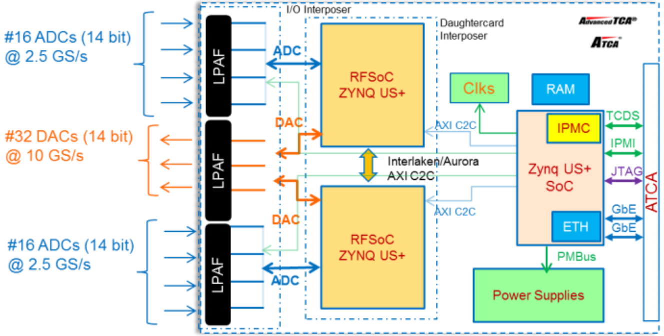

New LPAF (Low Profile Open-Pin-Field Array) connector: High-Speed High-Density analog & digital signals

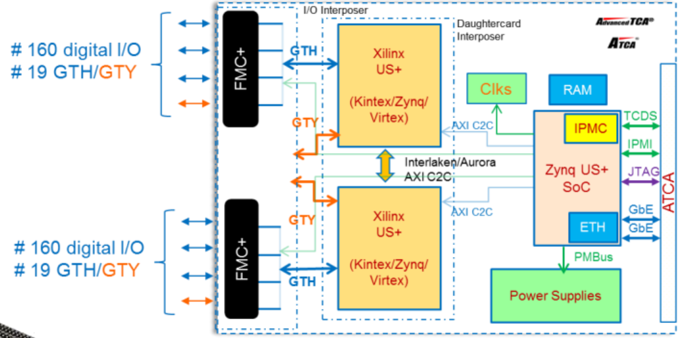

New FMC+: electrically & mechanically compatible with standard FMC

High-performance heterogeneous FPGA-GPU DAQ

- Modern photon science detectors generate huge raw data volumes (~120 Gb/s)
- Observed slow changes in synchrotron machine (e.g. current) → sec - hrs

PCIe readout card of UFO DAQ Platform

- Heterogeneous FPGA/GPU-based readout system, the UFO DAQ platform, has been developed
  http://ufo.kit.edu/ufo
- The core component is a “novel” Direct Memory Access (DMA) architecture
- Direct FPGA ↔ GPU communication enables real-time data processing

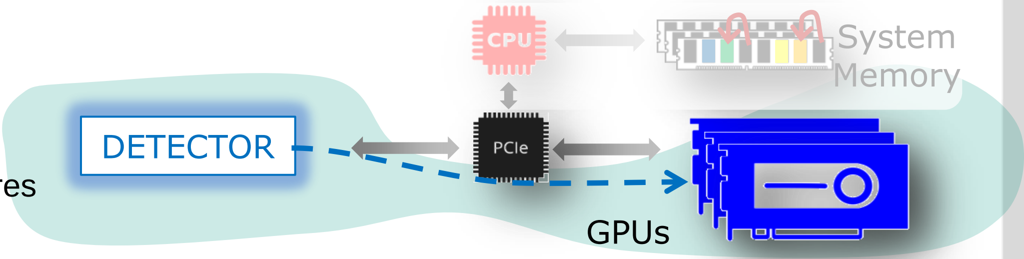

- DMA working close to theoretical limit of data link
- Data latency 5 times better than other DMA architectures

M. Caselle et al., JINST 12 C03015 (2017)

Multiple "GPU-Direct" architecture

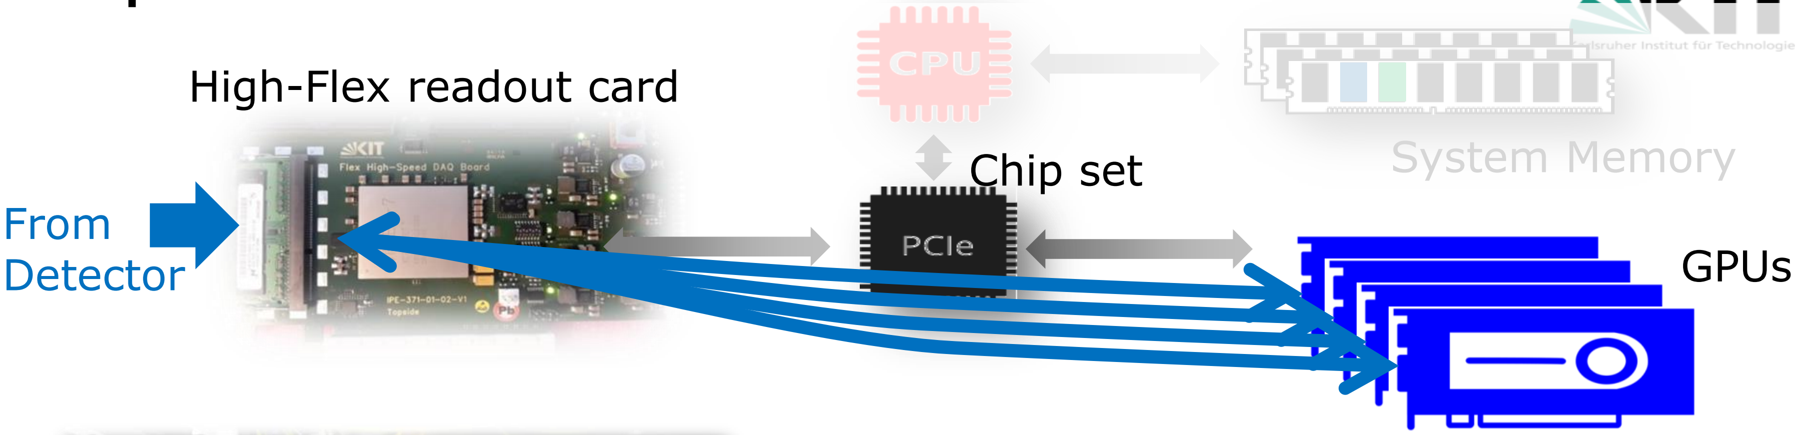

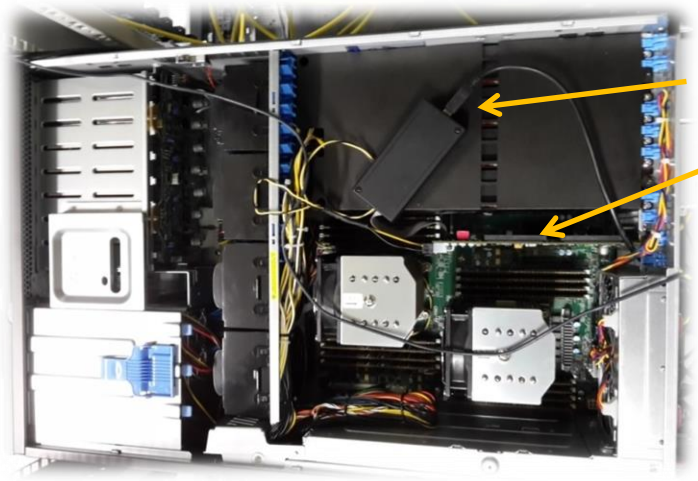

Only 20 ns are necessary to switch from GPU_1 to GPU_N

Ultra-fast computer tomography real-time 3D reconstruction (data rate 50Gb/s). DOI: 10.1109/TNS.2015.2425911. Real-time 2014 -Nara

Streaming Camera Platform for Scientific Applications. DOI: 10.1109/TNS.2013.2252528. Real-time 2012 – Lawrence Berkeley Laboratory

High-performance heterogeneous FPGA-GPU

Ultrafast X-ray Computer Tomography @ KARA (KIT)

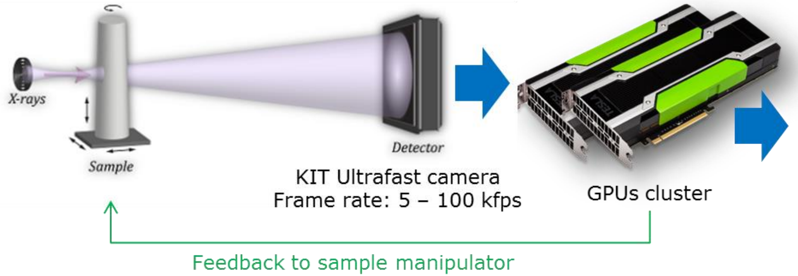

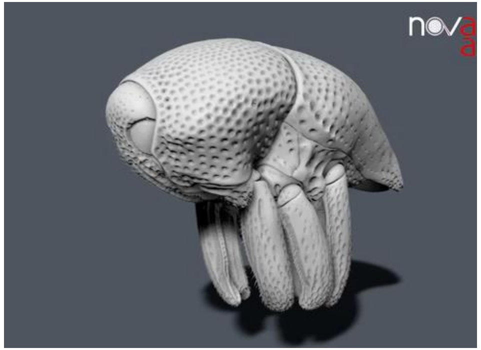

Reconstruction and semi-automatic segmentation by GPUs

3D spatial resolution: <1 µm + time resolution→ µs / ms

Data processing up to 50 Gb/s in real-time

M. Caselle et al., IEEE-RT DOI:10.1109/TNS.2013.2252528 (2013)

CMS low-level trigger system

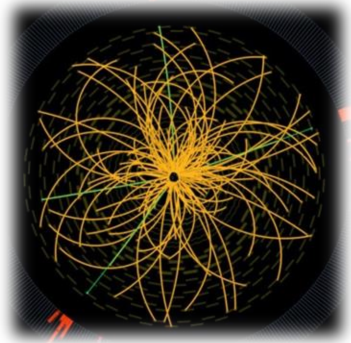

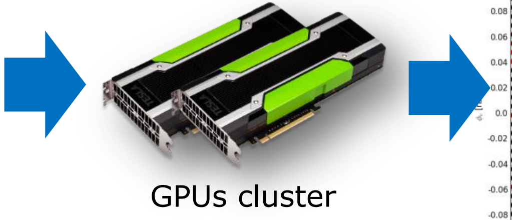

L1 trigger will require reconstruction of charged particles with transverse momentum > ~2 GeV/c

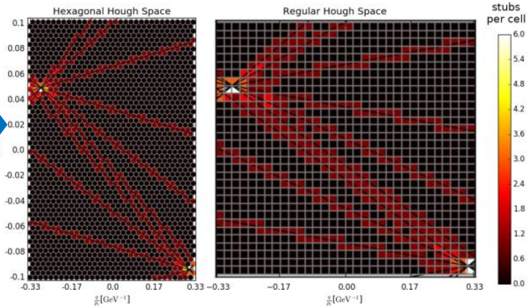

Hough space on GPU

CMS low-level trigger system based on FPGA-GPUs track reconstruction and fitting

Total data latency of 6.9 µs = 2 µs (data transfer) + 4.9 µs (GPU processing)

H. Mohr, M. Caselle et al., JINST 12 C04019 (2017)

Heterogeneous FPGA-GPU (Performance)

Data throughput KALYPSO → FPGA → GPU

- DMA working close to theoretical limit of data link (PCIe Gen3)
- Data throughput > 12.5 Gbyte/s both NVIDIA and AMD
- *High throughput → fundamental to sustain a continuous acquisition for long time (sec,- hrs)*

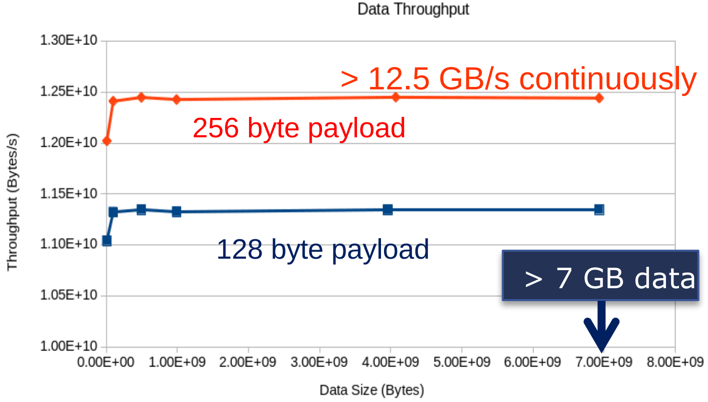

Round-trip time: FPGA →GPU →FPGA

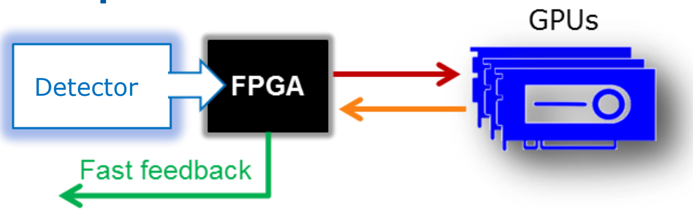

- NVIDIA (Tesla K40) : data latency < 2, jitter < 30 ns
- AMD (FirePro W9100): data latency < 1.4 µs, jitter 50 ns
- *Low-latency → fundamental for fast feedback to experimental system*

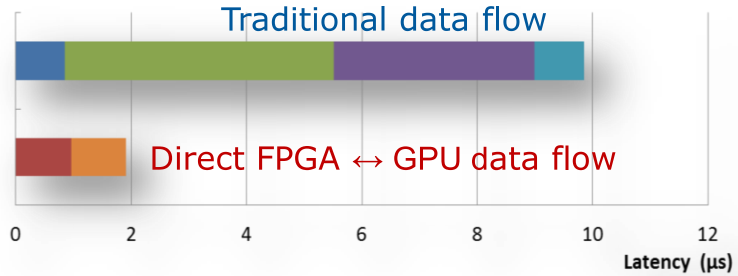

Versal – Adaptable Engines

Karlsruher Institut für Technologie

For traditionalists: This is the FPGA part

Some known facts

- 6 Input LUTs
- **Each CLB has 32 LUTs and 64 FF (4x density compared to US+)**
- 16 LUTs in a slice can be
  - a 64 bit RAM
  - 32-bit shift registers (SRL32) or two SRL16
- **Internal connection of LUTs possible**
- 4x clock, 4x set/reset, 16 clock enable
- **3 step voltage-scaling supported**

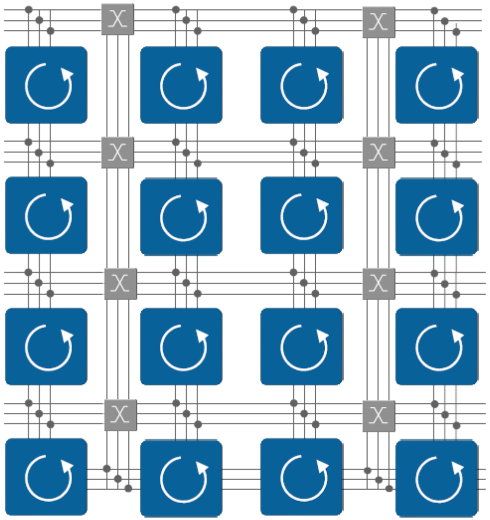

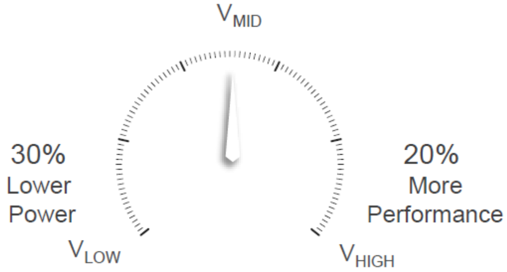

Versal – AI tile architecture

1.3 GHz VLIW / SIMD vector processors

Parallelity

- VLIW: 7+ operations / clock cycle
- SIMD: 512 bit vector datapath (8 / 16 / 32 bit & SPFP operands)
- Up to 128 INT8 MACs / clock cycle / core

Memory

- 16 KB Internal program memory
- 32 KB data memory (parallel)
- Integrated DMA logic

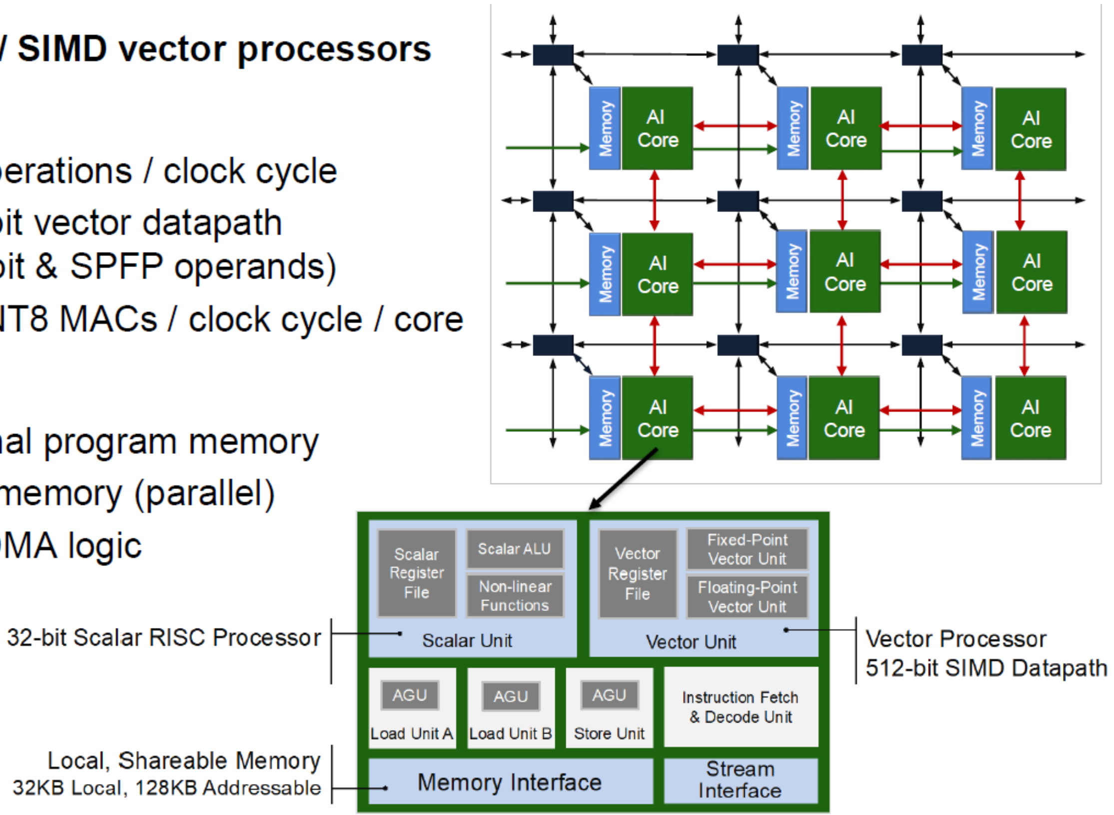

Advantage of Multi-Processor SoC technology

The user-friendly Linux applications is combined with the efficiency and throughput of a system fully implemented in the FPGA fabric

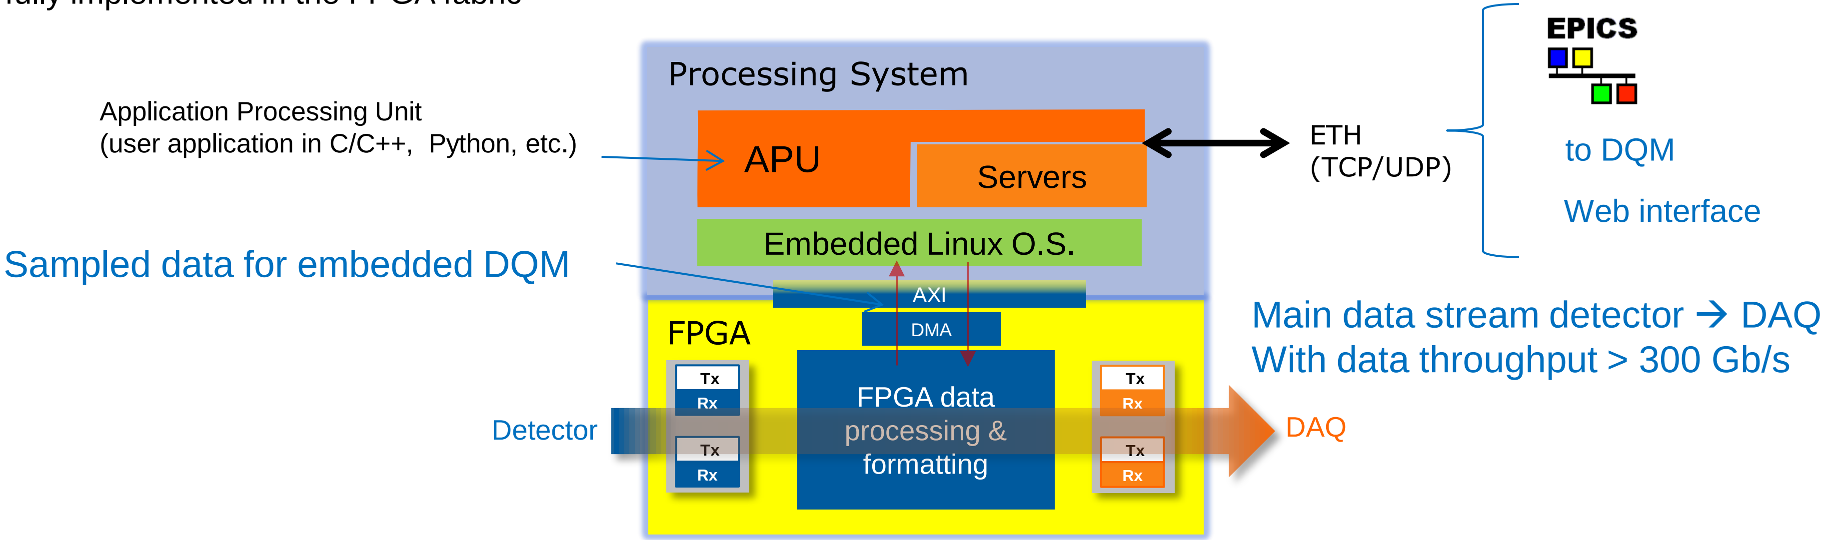

- Embedded ('slow-control') server on FPGA , i.e. EPICS server, etc.
- Detector ('calibration') routines running on FPGA
- High – granularity on-line monitoring of system and data quality ('DQM') on FPGA
- etc.

Advantage of Multi-Processor SoC technology

The user-friendly Linux applications is combined with the efficiency and throughput of a system fully implemented in the FPGA fabric

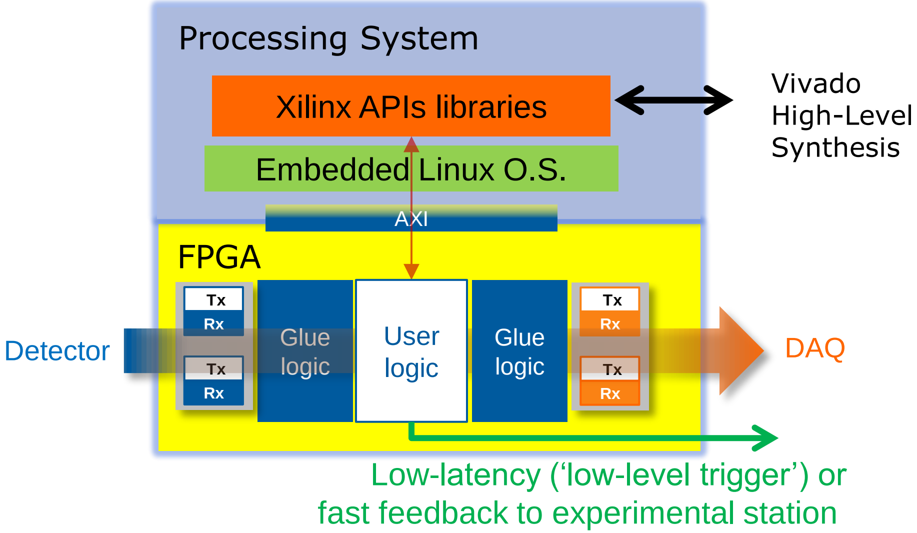

- PYNQ, reVISION: are an open-source project from Xilinx that makes it easy to programmable hardware without having to use HDL (verilog/VHDL) language.
- Low-level trigger based on Artificial Intelligent