


# Towards heterogeneous FPGA architectures and application examples

Oliver Sander



# Overview

- Part 1 -

FPGAs  
History and Status



- Part 2 -

MPSoC and selected  
application examples



- Part 3 -

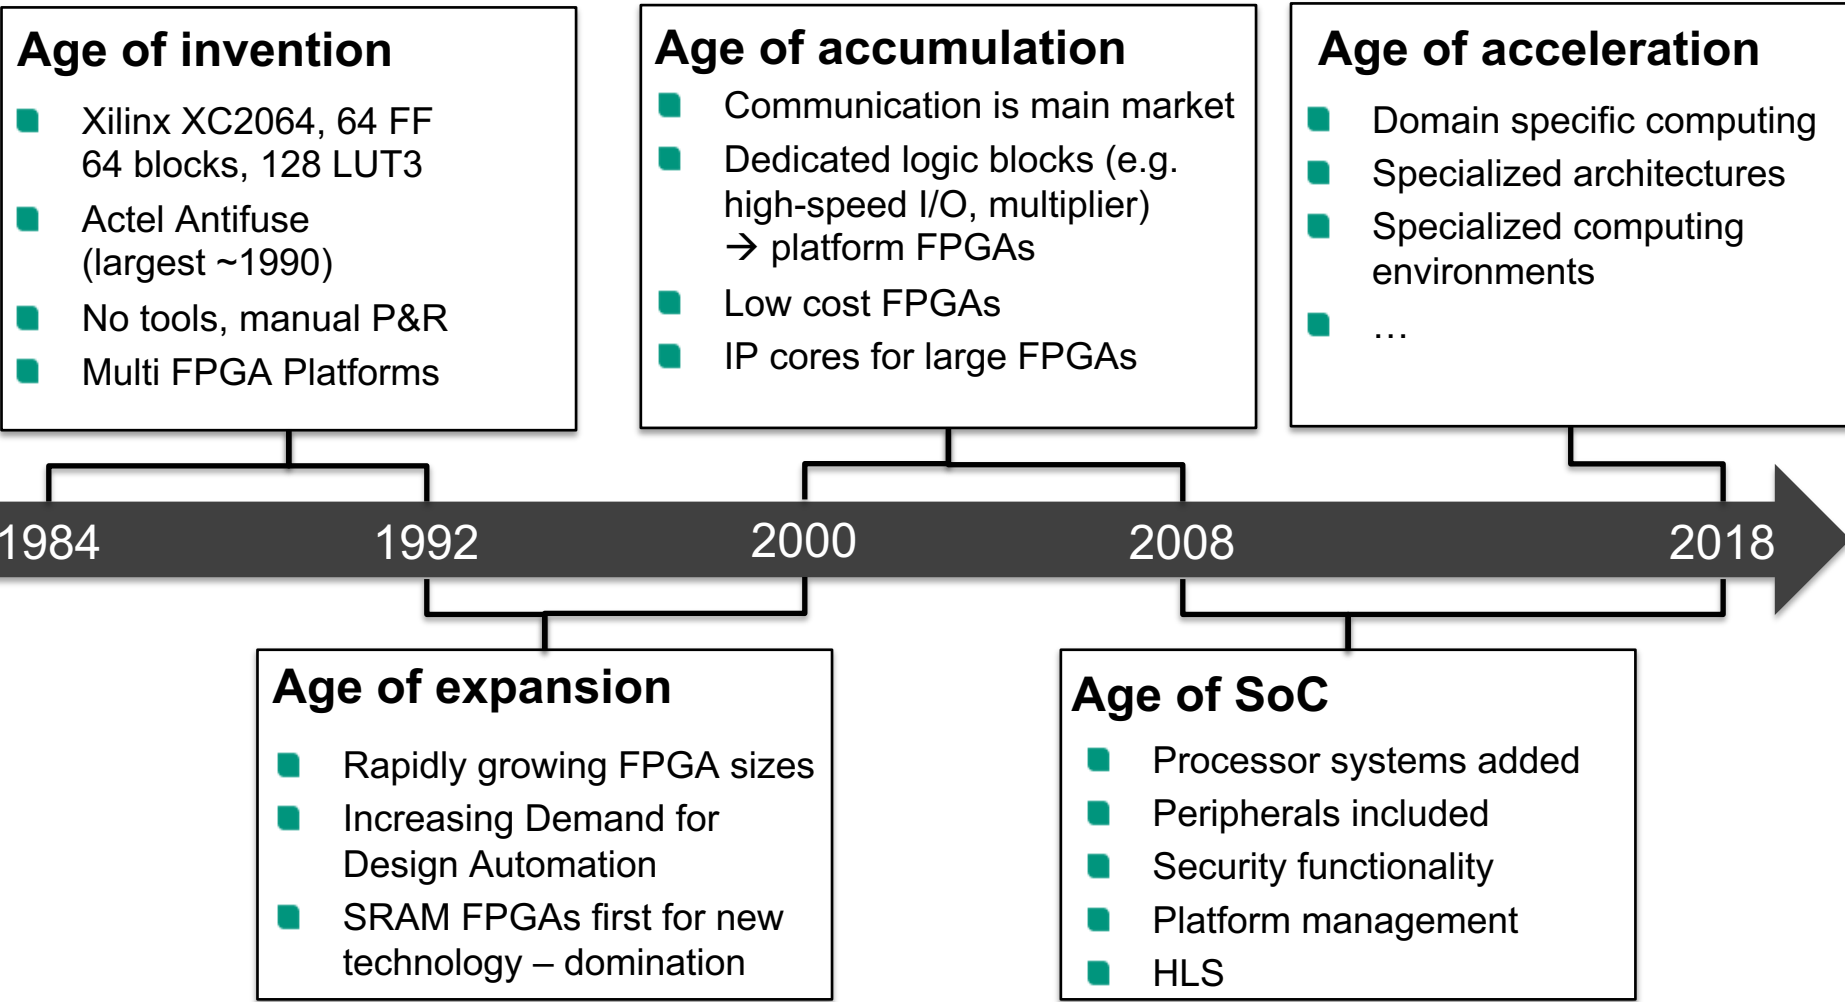
Next generation  
architecture



## Comments/Disclaimer

- (1) Talk content is biased towards Xilinx FPGAs. This is neither a statement nor a recommendation.
- (2) Talk focuses on high-end architectures to show technical development.
- (3) Content is a personal selection and not exhaustive.
- (4) Versal Information comes from XDF Frankfurt (links need to be added)

# Evolution of FPGAs



Ages are defined in a paper by Stephen Trimberger (2015) and in an interview with Ivo Bolsens (Xilinx CTO) in 2018

# FPGA market or which vendors did survive?

**achronix**  
SEMICONDUCTOR CORPORATION

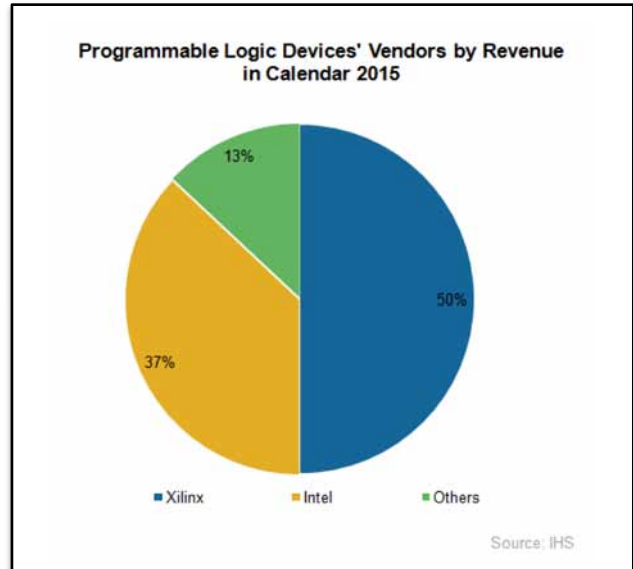
FPGA IP core for SoC designs

**intel** FPGA

SRAM based FPGAs  
Broad range

**LATTICE**  
SEMICONDUCTOR

Low power & cost efficient FPGAs



**MICROCHIP**


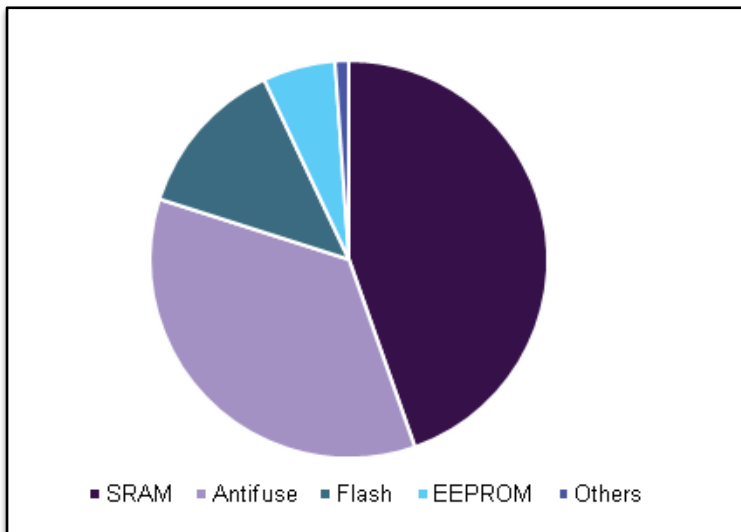
2 small SRAM FPGA families

**Microsemi**











Flash & Antifuse FPGAs

**XILINX**  
ALL PROGRAMMABLE™

SRAM based FPGAs  
Broad range




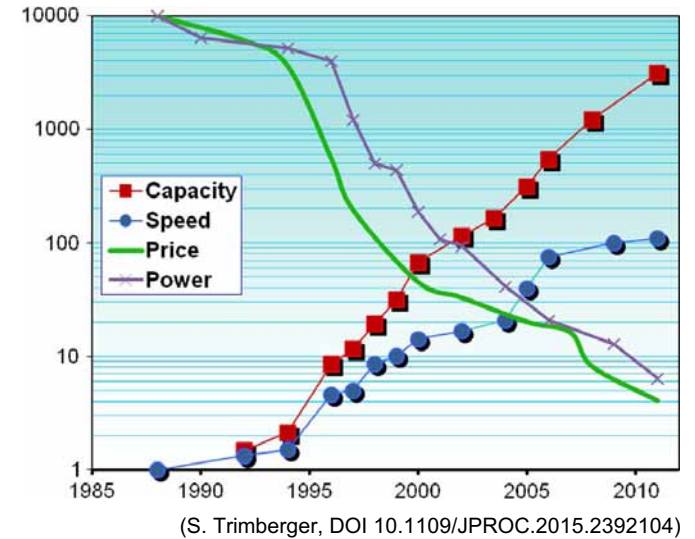


# What about Intel vs. Xilinx?

|   |  |   |   |  |
|---|--|---|---|--|
| <p>FPGAs for cost sensitive or mid-range products</p> | <ul style="list-style-type: none"> <li>• focus on logic/money</li> <li>• limited IO bandwidth</li> </ul>   |    |  |   |
| <p>High performance FPGAs</p>                         | <ul style="list-style-type: none"> <li>• maximum LUTs</li> <li>• maximum DSP slices</li> <li>• maximum internal memory</li> <li>• maximum IO bandwidth (30 Gbps, 58 Gbps)</li> </ul> |    |   |   |
| <p>FPGA + HBM</p>                                     | <ul style="list-style-type: none"> <li>• derived from HP FPGAs</li> <li>• integration of large memories (GB)</li> </ul>  |    |   |   |
| <p>FPGA + Processor System</p>                        | <ul style="list-style-type: none"> <li>• derived from HP FPGAs</li> <li>• multiple processors</li> <li>• memory and caches</li> <li>• peripherals</li> </ul>                         |   |   |  |
| <p>FPGA+Processor+ ADC/DAC</p>                        | <ul style="list-style-type: none"> <li>• derived from previous</li> <li>• high performance ADC/DAC</li> </ul>  |  |   |  |

# FPGA complexity over the years in numbers

- Limited feature requirements (transistors, wires) of SRAM FPGAs allowed early adoption of new technology nodes  
→ front-runner
- Exponential progress in compute power, memory, and bandwidth
- Dramatic increase in power efficiency
- Dramatic decrease of price per logic gate




(Dissertation C. Amstutz, 2016)

# Features in modern FPGA architectures

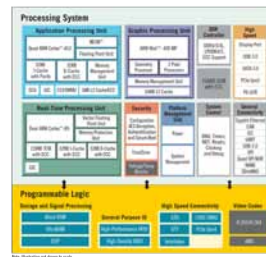
|  | Kintex UltraScale FPGA | Kintex UltraScale+ FPGA | Virtex UltraScale FPGA | Virtex UltraScale+ FPGA | Zynq UltraScale+ MPSoC | Zynq UltraScale+ RFSoc |
|--|------------------------|-------------------------|------------------------|-------------------------|------------------------|------------------------|
| MPSoC Processing System                    |                        |                         |                        |                         | ✓                      | ✓                      |
| RF-ADC/DAC                                 |                        |                         |                        |                         |                        | ✓                      |
| SD-FEC                                     |                        |                         |                        |                         |                        | ✓                      |
| System Logic Cells (K)                     | 318–1,451              | 356–1,143               | 783–5,541              | 862–3,780               | 103–1,143              | 678–930                |
| Block Memory (Mb)                          | 12.7–75.9              | 12.7–34.6               | 44.3–132.9             | 23.6–94.5               | 4.5–34.6               | 27.8–38.0              |
| UltraRAM (Mb)                              |                        | 0–36                    |                        | 90–360                  | 0–36                   | 13.5–22.5              |
| HBM DRAM (GB)                              |                        |                         |                        | 0–8                     |                        |                        |
| DSP (Slices)                               | 768–5,520              | 1,368–3,528             | 600–2,880              | 2,280–12,288            | 240–3,528              | 3,145–4,272            |
| DSP Performance (GMAC/s)                   | 8,180                  | 6,287                   | 4,268                  | 21,897                  | 6,287                  | 7,613                  |
| Transceivers                               | 12–64                  | 16–76                   | 36–120                 | 32–128                  | 0–72                   | 8–16                   |
| Max. Transceiver Speed (Gb/s)              | 16.3                   | 32.75                   | 30.5                   | 58.0                    | 32.75                  | 32.75                  |
| Max. Serial Bandwidth (full duplex) (Gb/s) | 2,086                  | 3,268                   | 5,616                  | 8,384                   | 3,268                  | 1,048                  |
| Memory Interface Performance (Mb/s)        | 2,400                  | 2,666                   | 2,400                  | 2,666                   | 2,666                  | 2,666                  |
| I/O Pins                                   | 312–832                | 280–668                 | 338–1,456              | 208–832                 | 82–668                 | 280–408                |

## New Features in Virtex Ultrascale+ (16 nm FinFET+)

- Ultra RAM Memory blocks (4kx72)
- Up to 8 GB HBM integrated DRAM (460 GB/s)
- 58 Gb/s PAM4 transceivers, 32 Gb/s
- PCI GEN3 (6x) and GEN4 (4x)
- 100G ethernet MAC with KR4-FEC & 150 G Interlaken cores

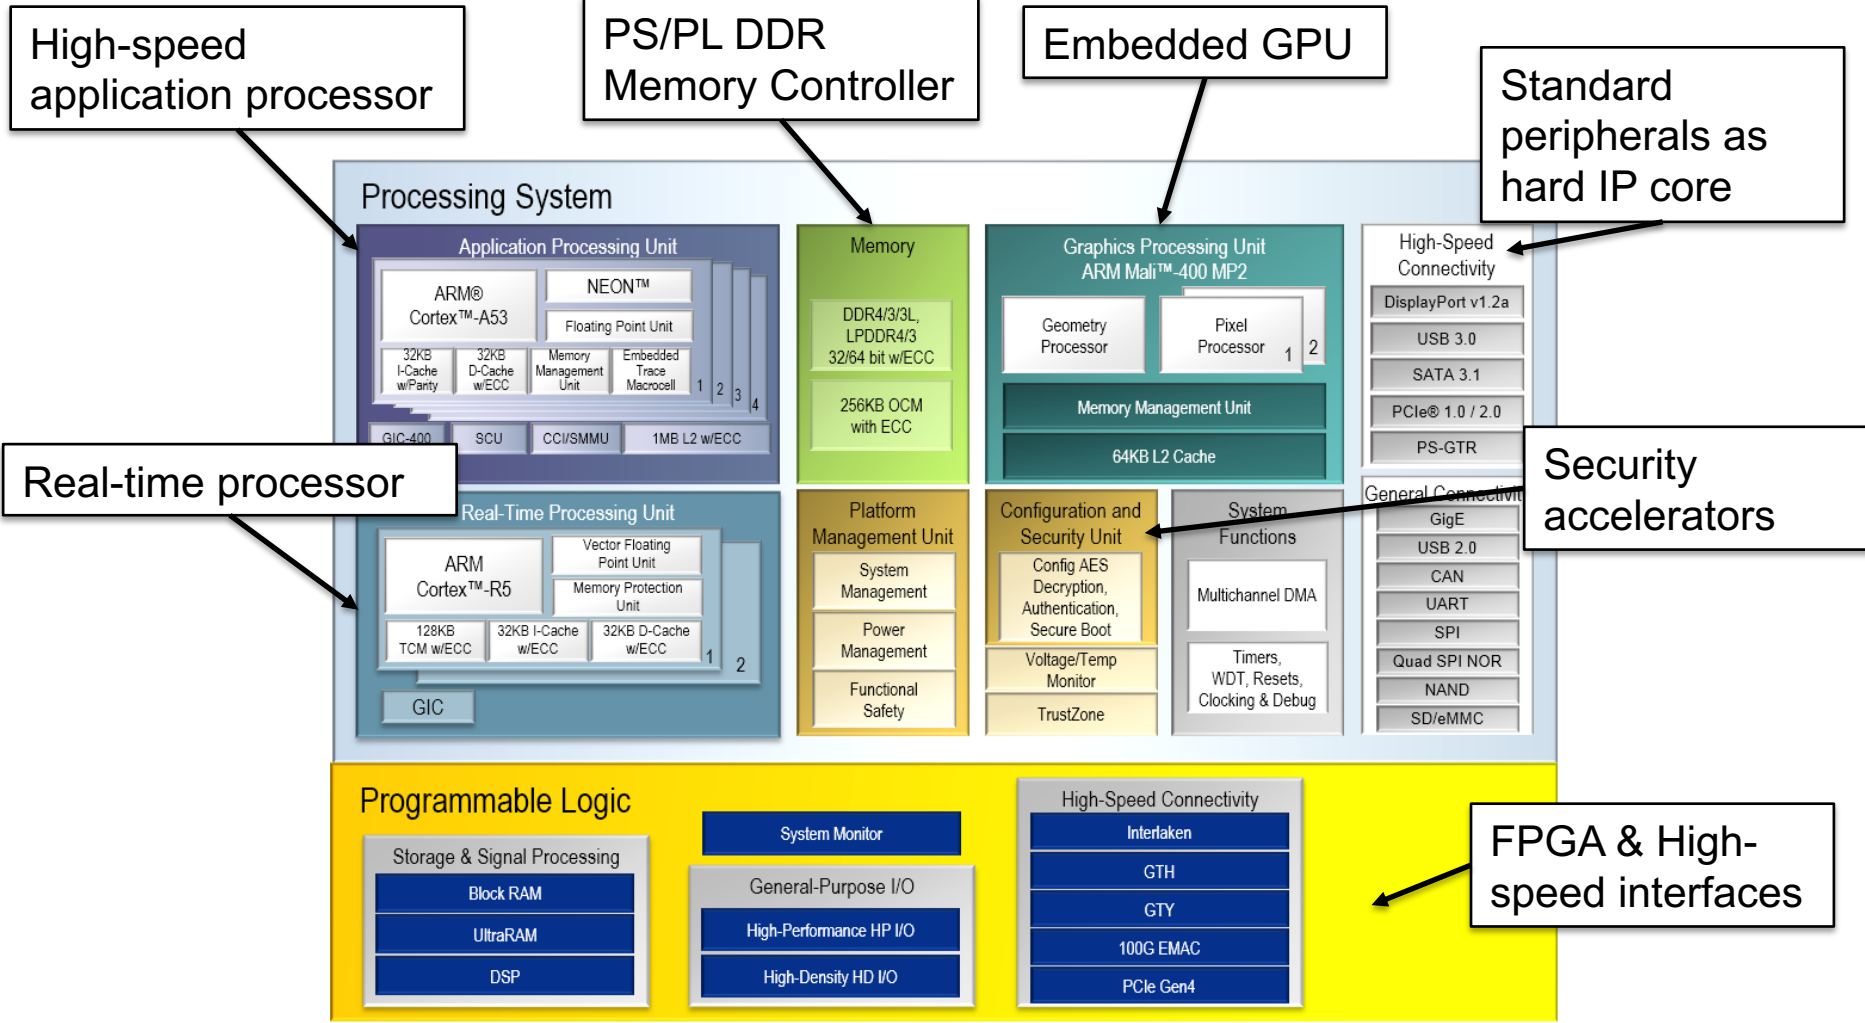


# - Part 2 - MPSoC and application examples



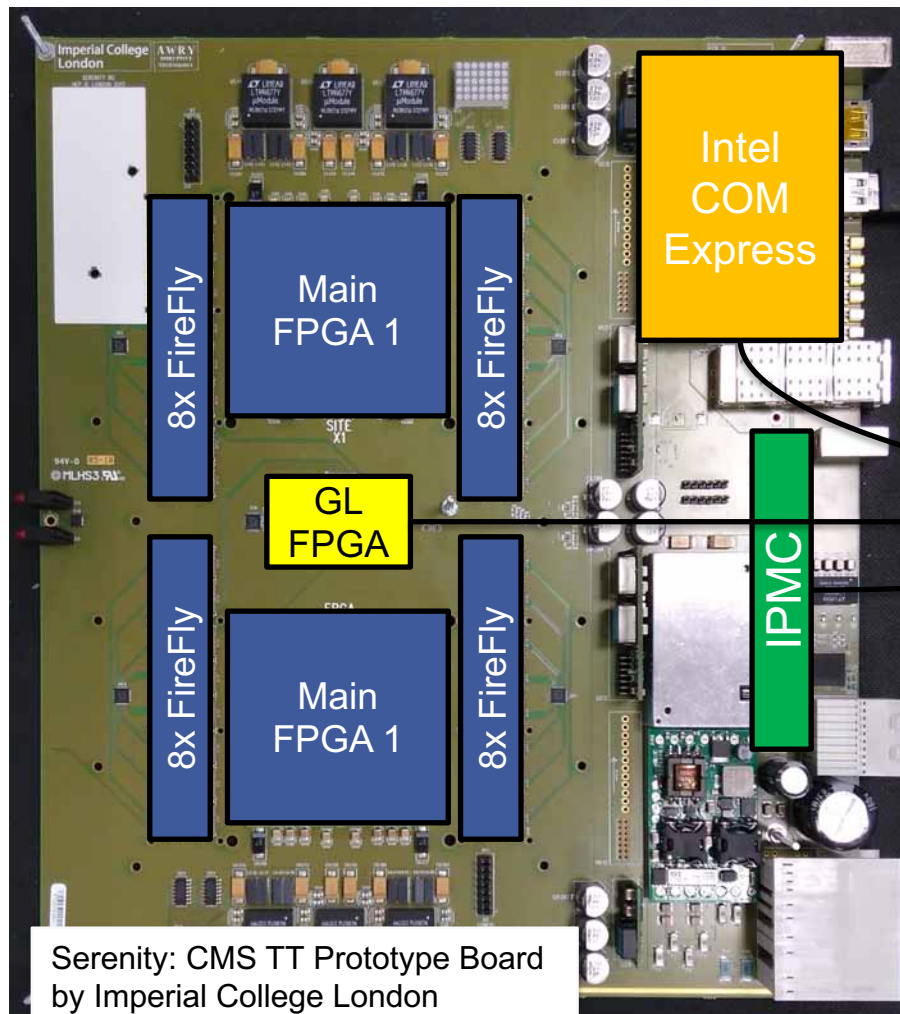


# Programmable MPSoC – Zynq Ultrascale+



<http://www.xilinx.com/products/silicon-devices/soc/index.htm>


# Example: Integrated IPMC for HL-LHC CMS L1 Track Trigger



- Intel Atom: High Performance interface to main FPGAs for dev, test, and calibration
- IPMC: Low-level platform management
- GL FPGA: interface conversion, glue logic



Why not integrate into one Zynq US+?



# Zynq US+ on multi purpose platform HiFlex 2

## Photon science

- New generation of detectors for beam diagnostics
- Diagnostics and stabilization of laser systems

## Superconducting sensors and quantum technologies


- Readout of superconducting sensor arrays
- Control- and readout of qubits

## 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



# Example: DAQ for the ECHo experiment

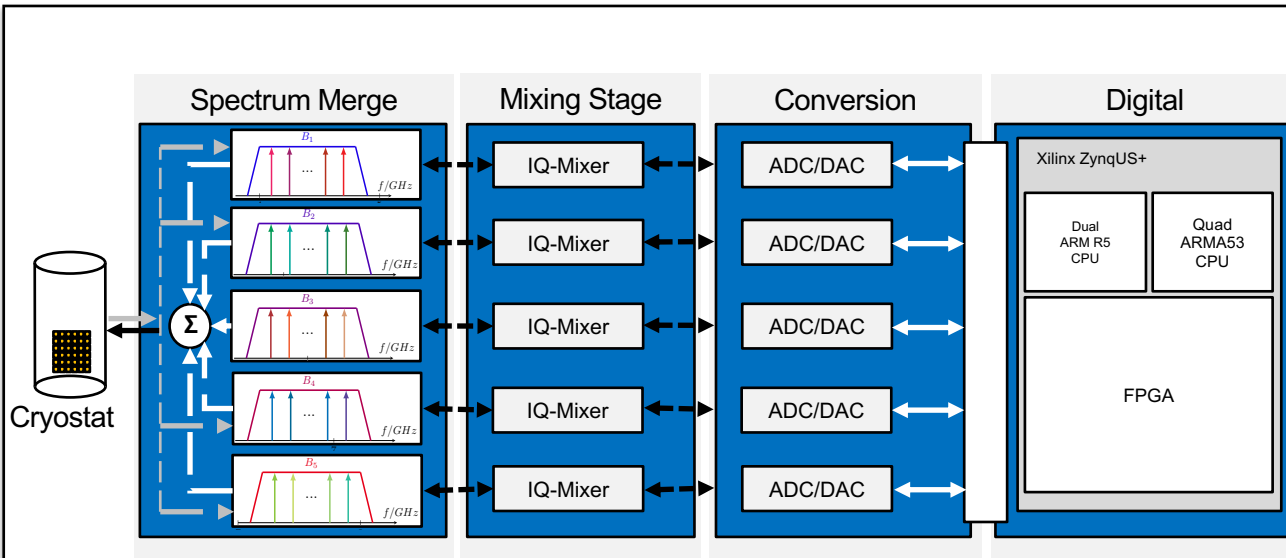


The **Electron Capture  $^{163}\text{Ho}$**  experiment **ECHo**<sup>[1]</sup> will measure the electron neutrino mass by analyzing the energy spectrum in the electron capture process of  $^{163}\text{Ho}$ .

## Technology

15x

- 800 superconducting sensors (MMC)
- 10 events per pixel per second
- 400 channels, one transmission line
- Frequency division multiplexing
- 4-8 GHz, one channel each 10 MHz



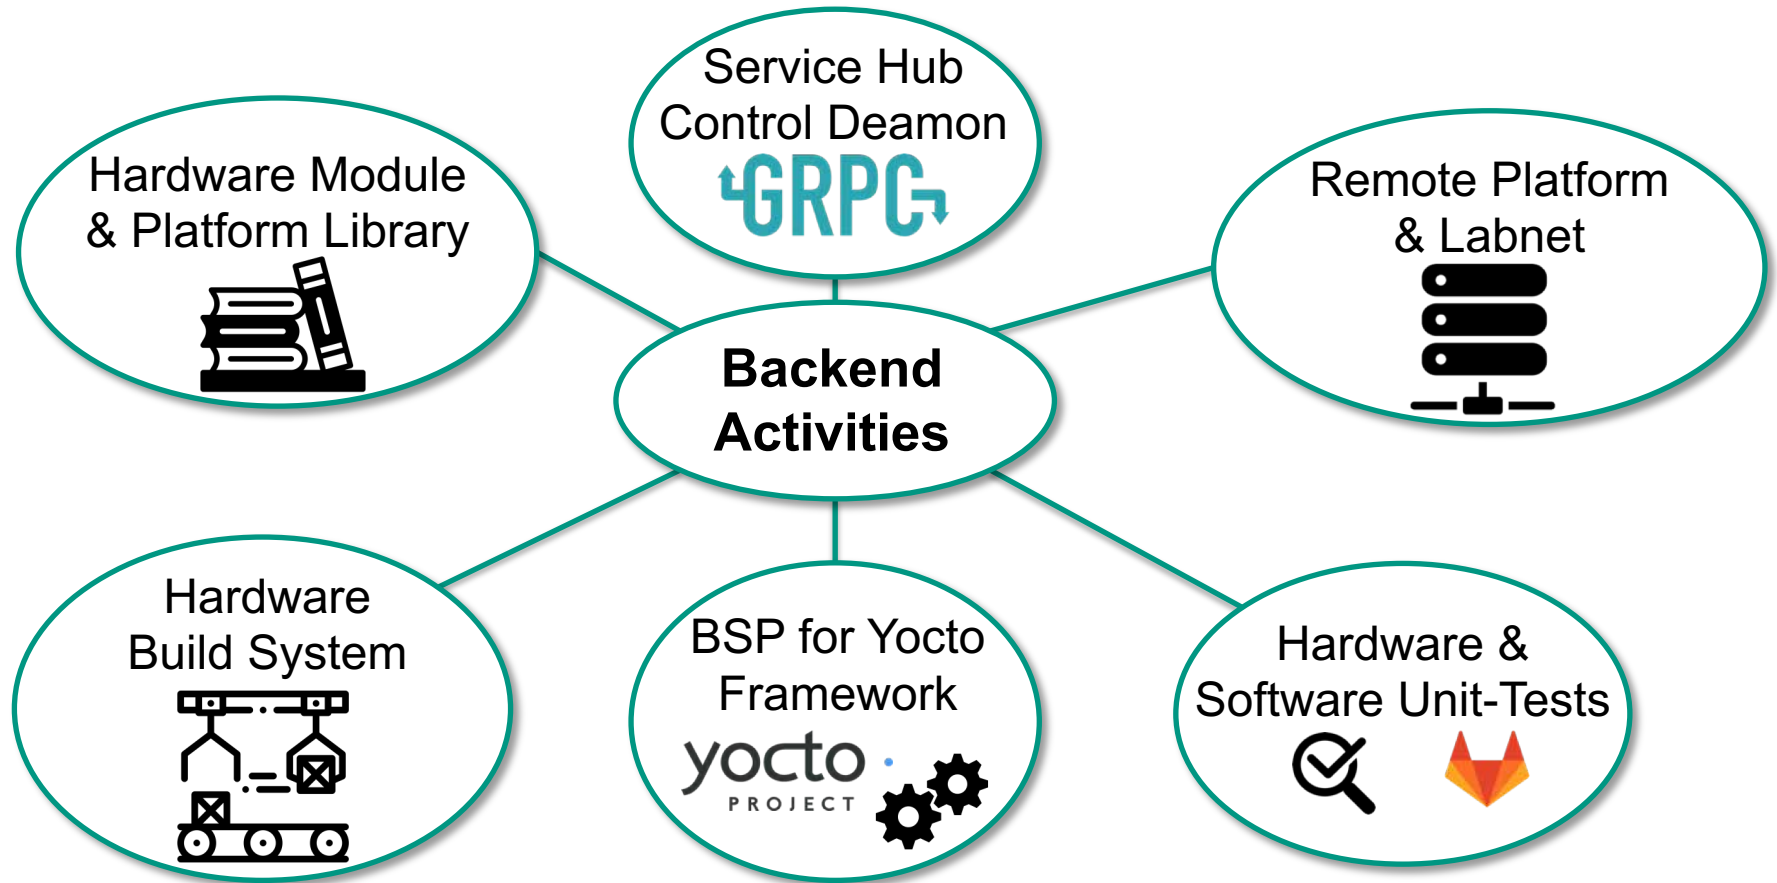
Software-defined Radio (SDR) system architecture

## MPSoC / FPGA

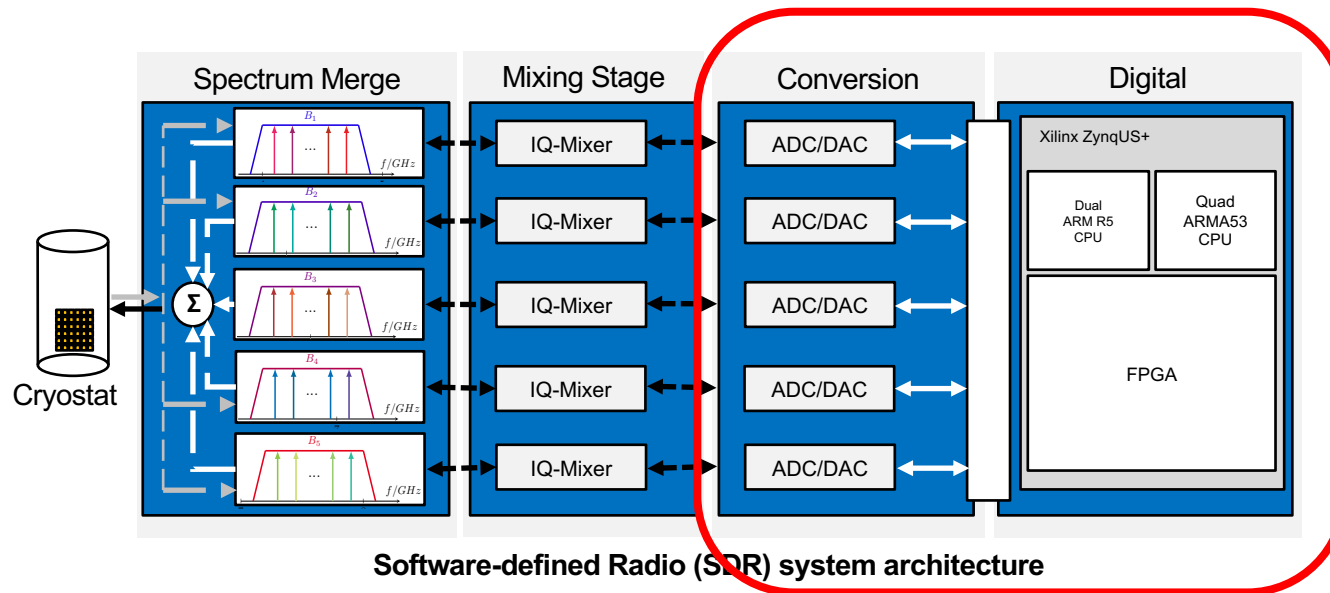
15x

- 160 Gbps Input
- 160 Gbps Output
- < 10 Mb/s to back-end storage server
- Full event processing on FPGA required
- Complex Calibration on processors with FPGA support

# IPE tooling environment for Zynq Ultrascale+



# Once more: DAQ for the ECHO experiment



Why not integrate the ADCs/DACs into a heterogeneous MPSoC platform?

# Zynq Ultrascale+ becomes more heterogeneous

- Xilinx integrated high-performance ADC/DACs → RFSoc

|   |                                    | Baseband                                     | Wireless Radio |        | Backhaul, Remote-PHY | Phased Array Radar / Radio |
|---|------------------------------------|--|----------------|--------|----------------------|----------------------------|
|   |                                    | ZU21DR                                       | ZU25DR         | ZU27DR | ZU28DR               | ZU29DR                     |
| RF Data Converters<br>Soft Decision, FEC  | 12-bit, 4GSPS ADC                  | -  | 8              | 8      | 8                    | -                          |
|   | 12-bit, 2GSPS ADC                  | -  | -              | -      | -                    | 16                         |
|   | 14-bit, 6.4GSPS DAC                | -  | 8              | 8      | 8                    | 16                         |
|   | SD-FEC                             | 8  | -              | -      | 8                    | -                          |
| Processing System &<br>Programmable Logic | Application Processor Core         | Quad-core ARM Cortex-A53 MPCore up to 1.5GHz |                |        |                      |                            |
|   | Real-Time Processor Core           | Dual-core ARM Cortex-R5 MPCore up to 533MHz  |                |        |                      |                            |
|   | High Speed Connectivity            | DDR4-2600, PCIe Gen3 x16, 100G Ethernet      |                |        |                      |                            |
|   | Logic Density (System Logic Cells) | 930K   | 678K           | 930K   | 930K                 | 930K                       |
|   | DSP Slices                         | 4,272  | 3,145          | 4,272  | 4,272                | 4,272                      |
|   | 33G Transceivers                   | 16   | 8              | 16     | 16                   | 16                         |

| Gen 1 |       | Gen2  |       | Gen3 |      |      |
|-------|-------|-------|-------|------|------|------|
| ADC   | DAC   | ADC   | DAC   | ADC  | DAC  |      |
| 4.096 | 6.554 | 2.275 | 6.554 | 5.0  | 10.0 | GSPS |

# - Part 3 - Next generation FPGA(?) architecture





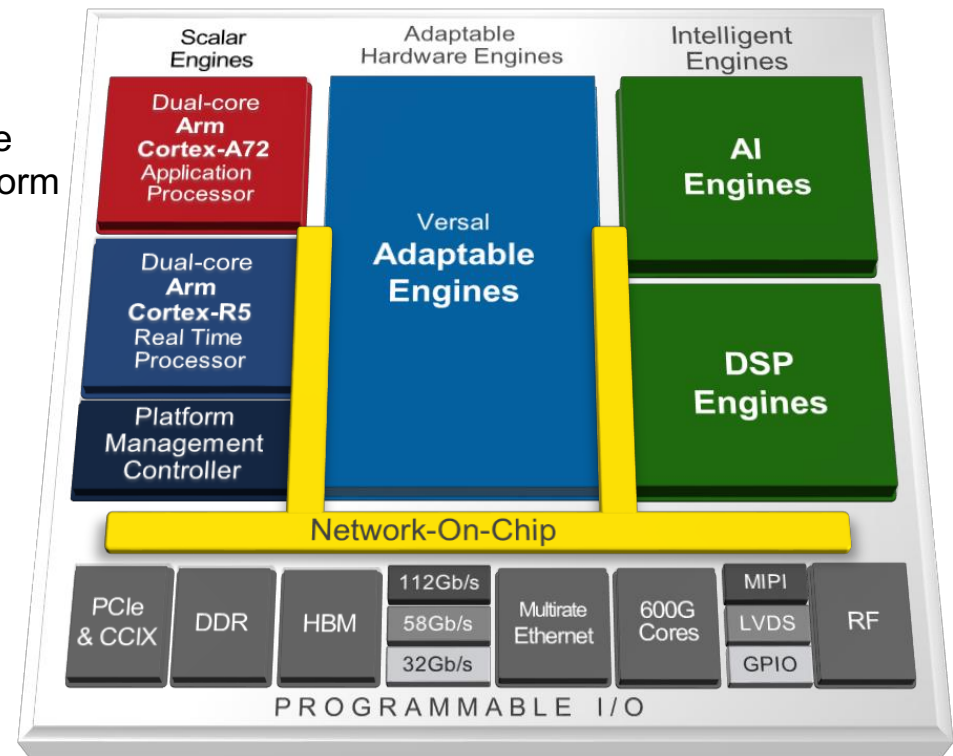
# Versal - Architecture Overview

## It is Xilinx' newest architecture

- More heterogeneous
  - More complex
- } Adaptive compute acceleration platform (ACAP)

## Key Features

- FPGAs + Processors + AI Engines
- Network on Chip backbone
  - High bandwidth & low latency
  - Guaranteed QoS
  - Memory mapped
  - built in arbitration
- Complex memory hierarchy (LUTRAM, BRAM, UltraRAM, Accelerator RAM, HBM, DDR)
- + optimizations in FPGA components

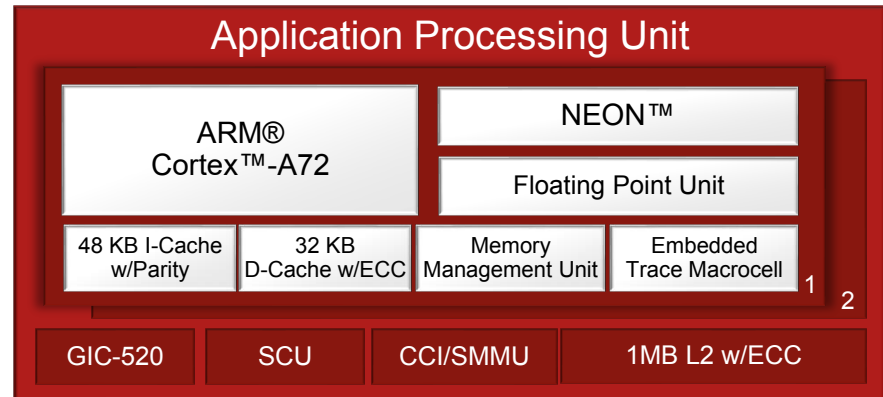


|  |           |           |
|--|-----------|-----------|
| System Logic Cells (K)                     | 540-1,968 | 352-2,154 |
| Hierarchical Memory (Mb)                   | 68-191    | 40-324    |
| DSP Engines                                | 928-1,968 | 472-3,984 |
| AI Engines                                 | 128-400   | -         |
| Processing System                          | ✓         | ✓         |
| Serial Transceivers (NRZ, PAM4)            | 8-44      | 12-66     |
| Max. Serial Bandwidth (full duplex) (Tb/s) | 2.9       | 4.2       |
| I/O  | 346-692   | 238-778   |
| Memory Controllers                         | 2-4       | 1-6       |

# Versal – Scalar Units

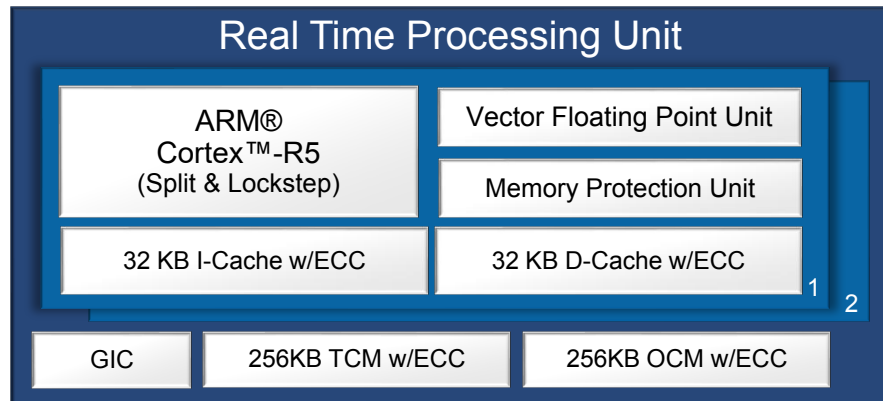
## Dual-Core ARM Cortex-A72 application processors

- Arm-v8A architecture
- Up to 1.7 GHz
- 2x single-threaded performance (DMIPS Versal vs. Zynq US+)



## Dual-Core ARM Cortex-R5 real-time processors

- Arm-v7R architecture
- Up to 750 MHz
- Low latency and deterministic
- Supports lock-step
- Internal memory



## Peripherals



- Ethernet, SPI, I2C, CAN, UART, GPIO, USB, timer-counter, watchdog

# Versal – Adaptable Engines

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 - DSP blocks

## New key features

- More than 1 GHz of performance
- Integrated FP32, FP16 floating point
- Integrated complex 18x18 operations
- SIMD support for add/sub/acc (dual 24 bit, quad 12 bit)



# Versal - AI tile architecture

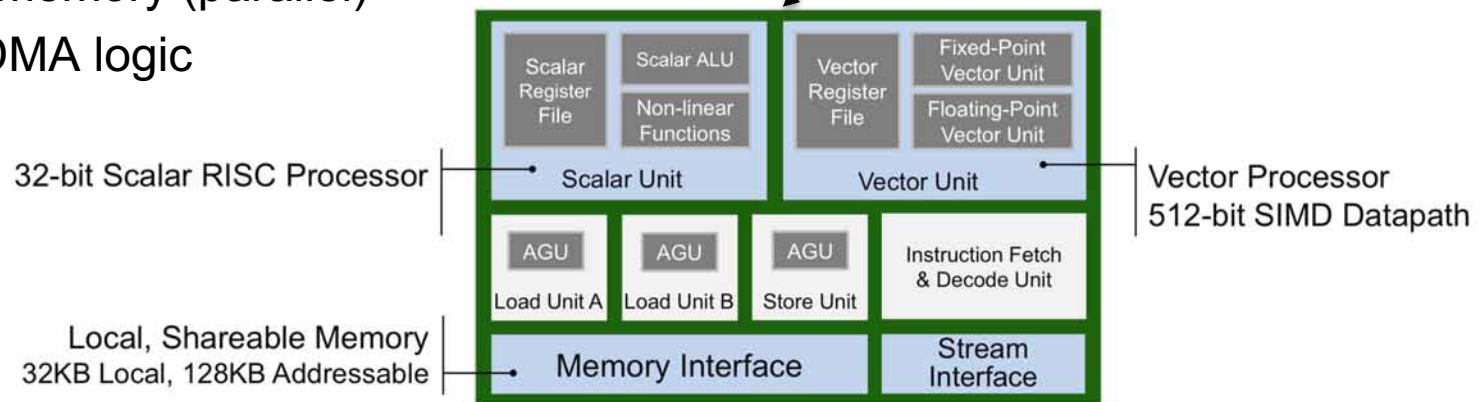
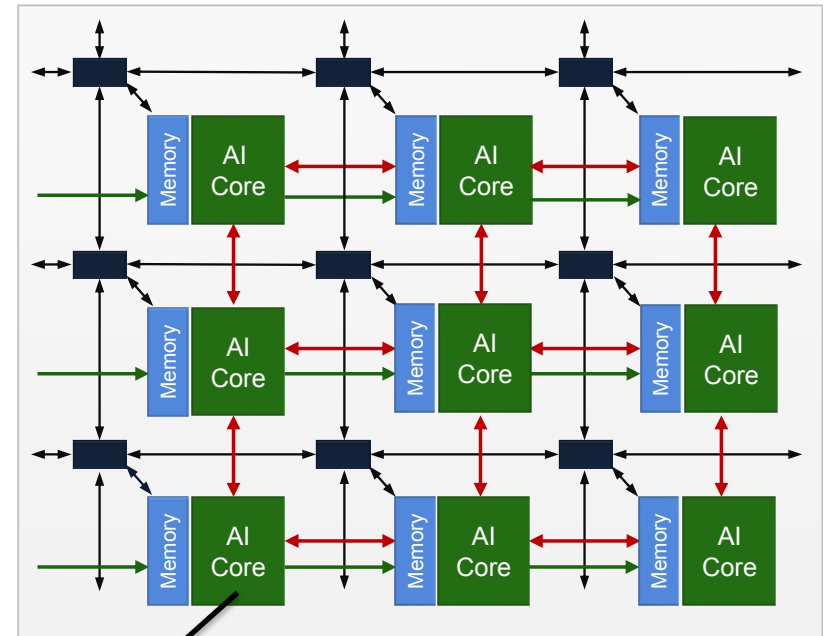
## 1.3 GHz VLIW / SIMD vector processors

### 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



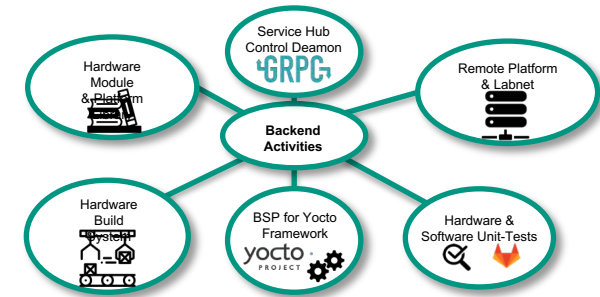
# Conclusion

- FPGAs become more and **more heterogeneous** devices
  - Zynq US+: FPGA & **CPU & Peripherals**
  - RFSoc: Zynq US+ & **ADC & DAC**
  - ACAP: FPGA & CPU & Per. & **VLIW/SIMD**



- Enables **high functional integration**  
(including control, calibration, and test software)

- Giant leaps in **tooling required** to leverage potential



- KIT IPE strongly believes in benefits of heterogeneous architectures → baseline for various projects



Thank you

Left intentionally blank