Flavor Physics Workshop Neckarzimmern

Moving from HEP to Industry with FPGAs

Christian Färber

Neckarzimmern - Germany

15. March 2023

Agenda

Moving from HEP to Industry with FPGAs

- Career Path
  - Heidelberg University
  - Vector Informatik
  - CERN
  - Thales
  - Intel
- What I found useful
  - Interviews
  - General Thoughts
  - Paths of FPGA Colleagues
- Summary
- FPGA Introduction
- FPGA compute acceleration

Field Programmable Gate Array

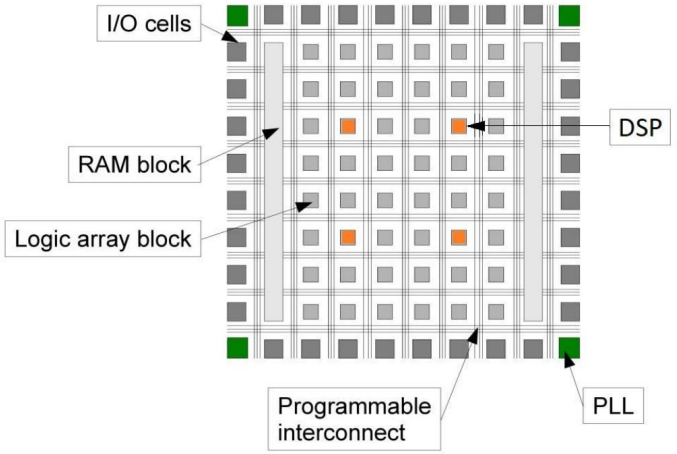

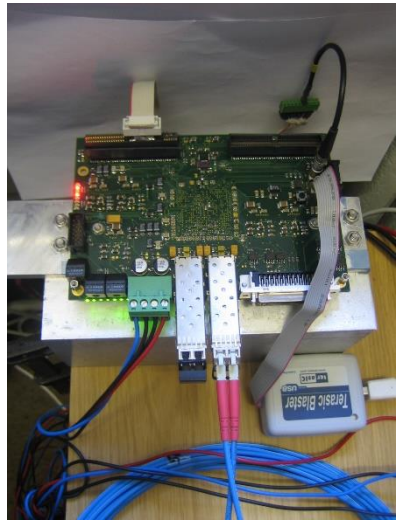

Source: https://www.intel.de/content/www/de/de/products/details/fpga/development-kits/stratix/10-dx.html

Section:
Career Path

Sub-Topics:

- Heidelberg University
- Vector Informatik
- CERN
- Thales
- Intel

2006 - 2009

Heidelberg University

Internship, Diploma

- Investigating Outer Tracker ageing
- Development of automatic ageing scanner
- Adding oxygen to counting gas and investigate ozone creation to reduce ageing

- Skills:
  - Undertaking irradiation tests in the lab
  - Implement a lab readout chain as also lab experiment control
  - Analyzing data with C++ and ROOT

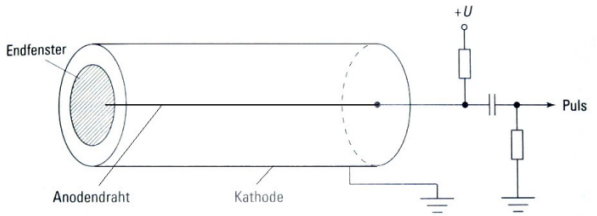

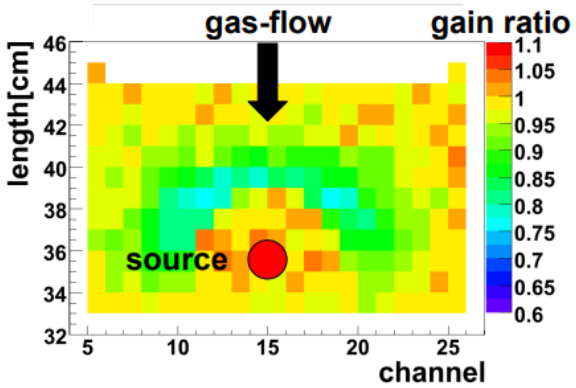

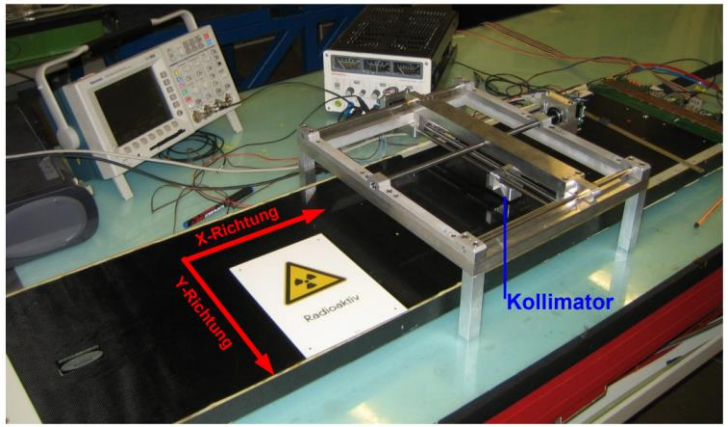

https://www.physi.uni-heidelberg.de/Publications/Diplomarbeit.pdf

2009 - 2014

Heidelberg University

PhD

- Prototyped FPGA-based readout board for potential OT upgrade
  - FPGA-based Time to digital converter (t_{res}<1ns)
  - FPGA irradiation tests at MPIK and HIT
  - Collaboration with engineers from PI workshop
- Support of Outer Tracker commissioning
- Outer Tracker gas monitoring
- Skills:
  - Electronics design with FPGAs, optical high-speed transceivers, VHDL, C++, analysis of experiments, presenting at conferences, ...
  - Developed muscle to work into complex new fields, having a high frustration tolerance and the mindset to get things really working
  - With HEP great preparation for joining academia and industry!

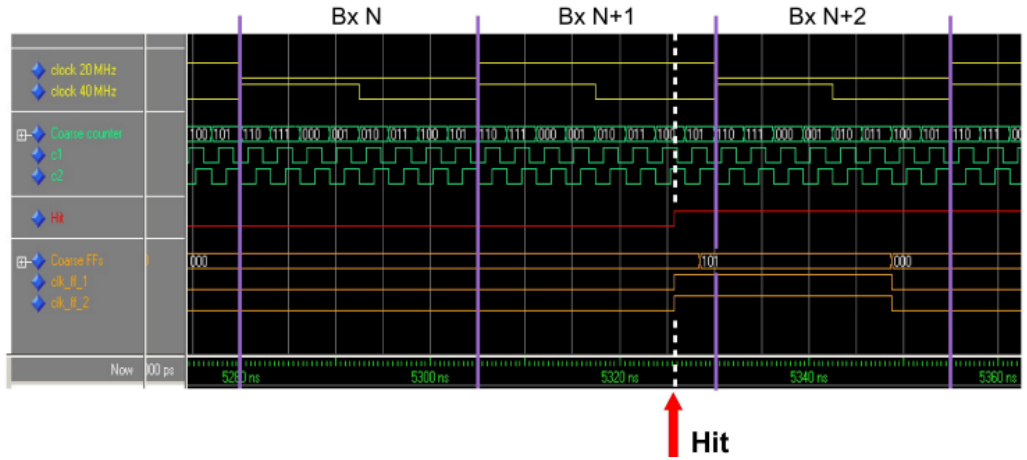

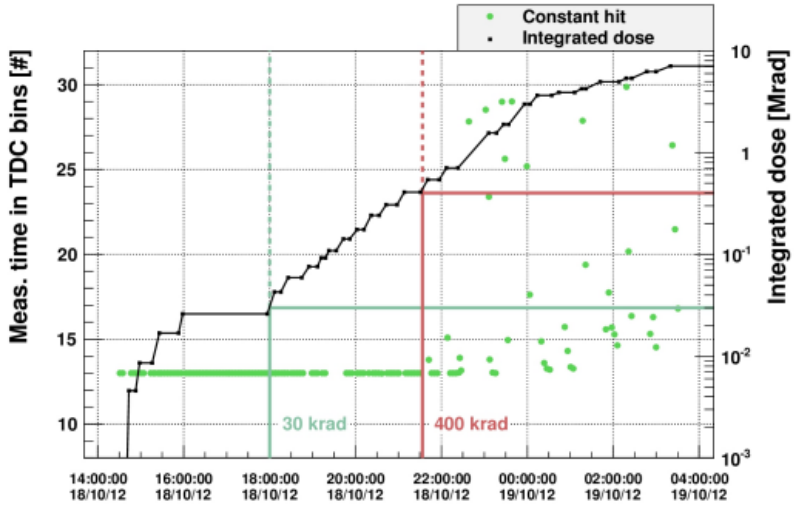

https://www.physi.uni-heidelberg.de/Publications/Dissertation_Christian_Faerber.pdf

Section:
Career Path

Sub-Topics:

- Heidelberg University
- Vector Informatik
- CERN
- Thales
- Intel

Vector Informatik

Hardware Development Engineer

- Vector: Develops software tools and components for networking of electronic systems based on serial bus systems for automotive
- Project: FPGA development of a high-speed interface and a memory synchronization block for the VX1000 hardware, used for vehicle ECU calibration
- Team of 5 FW engineers
  - Next to software team and PCB team, in total 20
- Programming Guide for VHDL & C++
- Industry timelines
- Working time
- Trainings necessary for CAN …

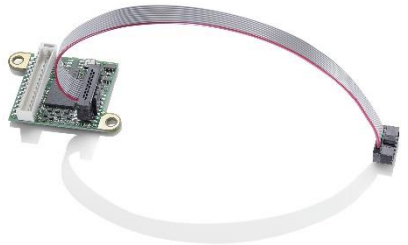

Readout chain

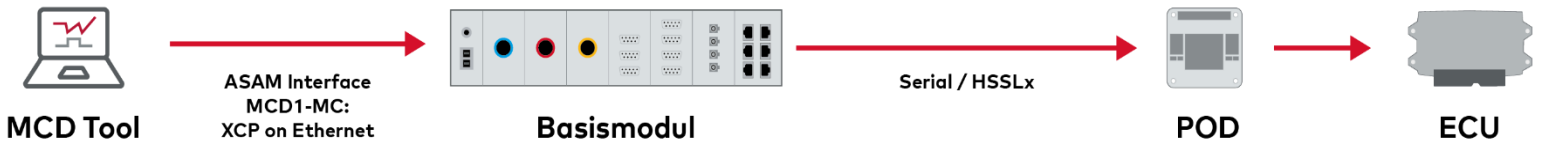

Source: https://www.vector.com/us/en/

Vector Informatik

How I got this job

- Job portals / social networks
  - Stepstone
  - For large companies check their career portal!
  - LinkedIn / (in the past Xing)
- Job-Messe University Heidelberg
- Created a professional resume – 2 pages only!
- Applied for several positions and companies
- Shared and discussed experience with my FPGA network
  - Dry run of interview with people from the field

Section:
Career Path

Sub-Topics:

- Heidelberg University
- Vector Informatik
- CERN
- Thales
- Intel

2015 - 2018

CERN - Openlab

Senior Applied Fellow

- HTCC - CERN Openlab collaboration with Intel
- Feasibility study to use FPGA-based compute acceleration in the HPC sector of High-Energy physics
  - RICH photon reconstruction
  - Calorimeter raw data encoding
  - CNN inference study with MNIST
- Highlight: Prototypes of the Intel® Xeon®+FPGA hybrid server-processor (HARP - Intel Hardware Accelerator Research Program)
- Got to know Intel development team in Hillsboro very well

CERN - Openlab

RICH PID

- Calculate Cherenkov angle Θ_c for each track t and detection point D, not a typical FPGA algorithm
- RICH PID is not processed for every event, processing time is too long!
- 748 clock cycle long pipeline written in Verilog
  - Additional blocks developed: cube root, complex square root, rot. matrix, cross/scalar product,...
- Lengthy task in Verilog with all test benches

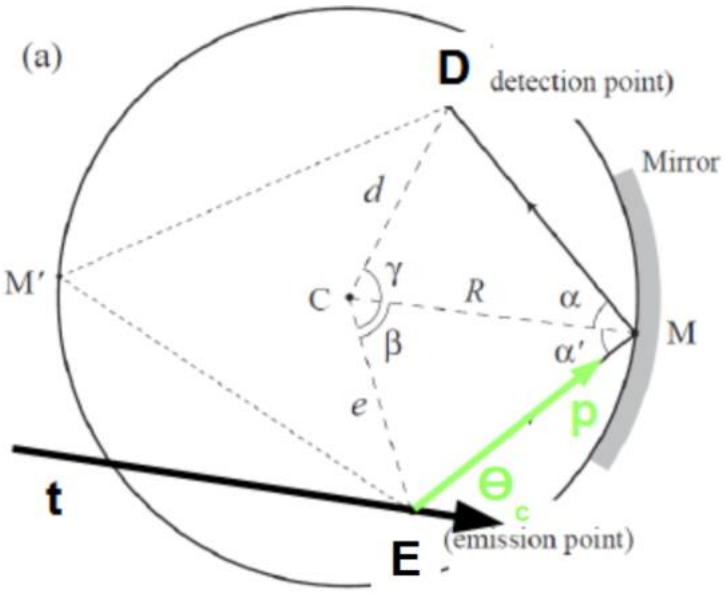

Reference: LHCb Note LHCb-98-040

2015 - 2018

CERN - Openlab

RICH PID

- Acceleration of up to factor 35 with Intel® Xeon®+FPGA
  - Theoretical limit of photon pipeline: a factor 64 with respect to single Intel® Xeon™ thread, for Arria™ 10 a factor ~ 300
  - Bottleneck: Data transfer bandwidth to FPGA, caching can improve this, tests ongoing
- Energy efficiency improvement
- Later implemented with OpenCL and standard PCIe acceleration cards

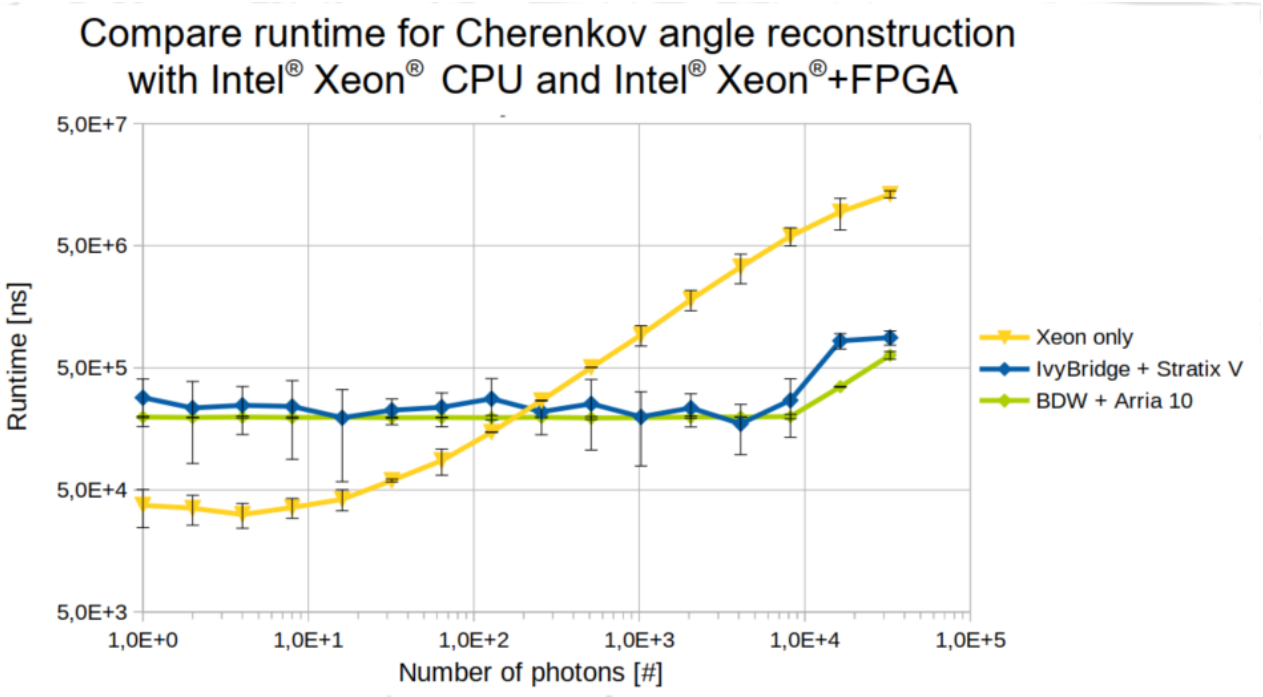

https://indico.cern.ch/event/669298/attachments/1551772/2433879/CERN_Computing_Seminar_Faerber_20171031.pdf

Section:
Career Path

Sub-Topics:

- Heidelberg University
- Vector Informatik
- CERN
- Thales
- Intel

2018 - 2019

THALES

Thales - MLS

Senior FPGA Design Engineer

- Thales Group is a French multinational company that designs, develops and manufactures electrical systems for the aerospace, defence, transportation and security sectors.
- Algorithm development and sensor optimization for applications in the rail industry
  - Got trainings in **CENELEC EN 50126, …**
- Fiber Bragg Gating based axle counting system
  - Simulation of sensor geometries
  - Rail field tests
  - Regular presentation of project status to Thales MLS management
  - Collaboration with product group and business unit supporting early customers
- Worked in a team of 6 FPGA design engineers, depart. 100.
- Direct manager and project manager

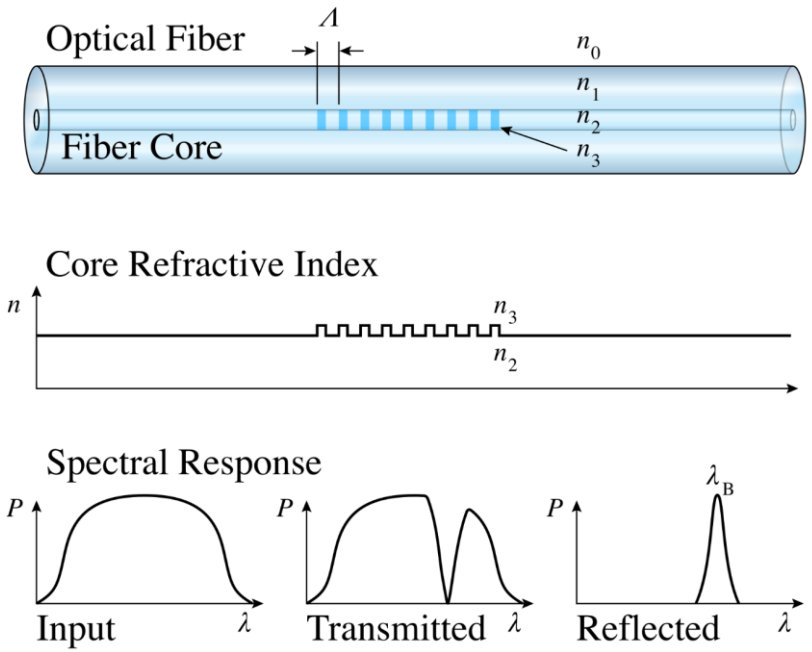

https://en.wikipedia.org/wiki/Fiber_Bragg_grating

Thales - MLS

Axle Counting

- Why axle counting?
- Old axle counting used magnet field
- New idea to measure bending of rail with FBG
- Replace expensive copper with cheaper fiber
- No power at sensor location needed O(km)

https://www.thalesgroup.com/en/countries/europe/germany/transportation/rail-field-equipment

Thales - MLS

Axle Counting

FORC

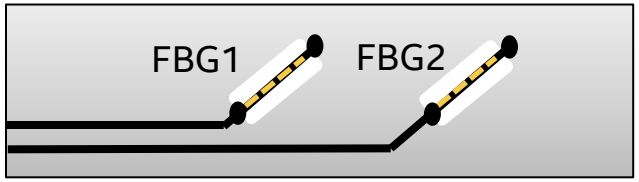

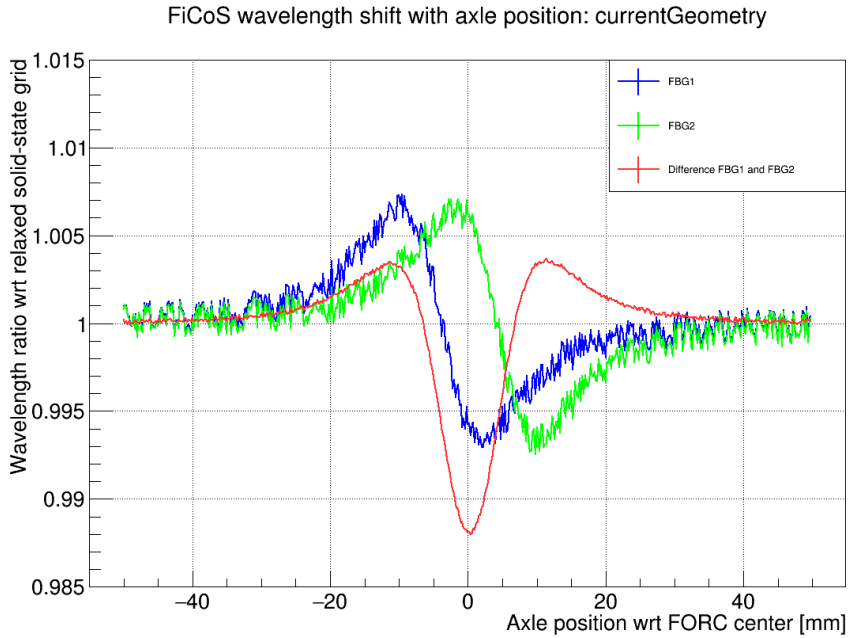

Working at Thales is a great possibility to influence modern transportation ☺

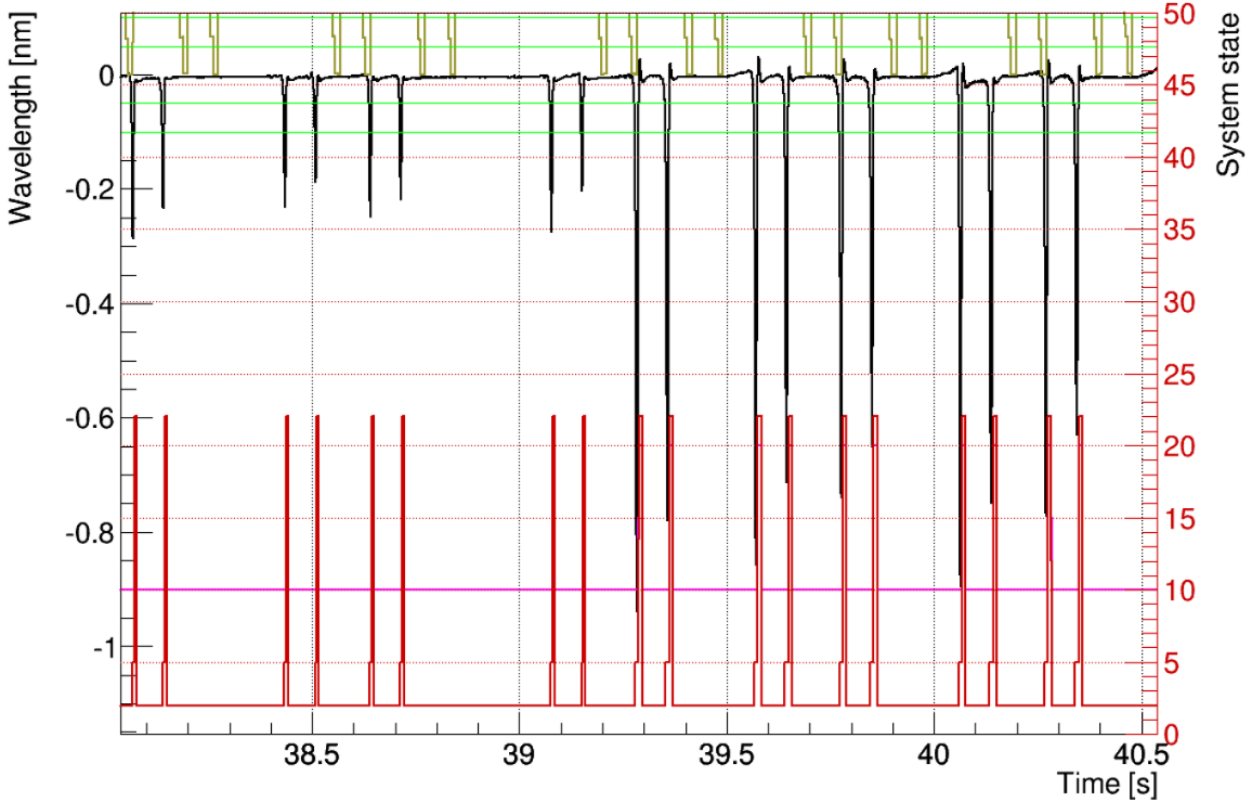

Great to see your algorithm working on the FPGA in the readout electronics for fields tests with real trains

Section:
Career Path

Sub-Topics:

- Heidelberg University
- Vector Informatik
- CERN
- Thales
- Intel

Since 2019

Intel - PSG

Senior Field Application Engineer

- General: Technical support of customers with freedom on which topics I work considering reaching Design Win
- How I got this job – my network
- No training, second day onside at customer
- Port of algorithms to FPGAs with goal to increase performance and reduce energy usage as also cost of future data centers
- Algorithms for large data center customers
  - Data compression
  - File parsing at line speed
  - Data analytics
- Not typical but possible:
  - Initiated that Intel participates in EU H2020 Daphne
  - Writing research papers and supervise students
  - Present research results at conferences, e.g.: HiPEAC2023

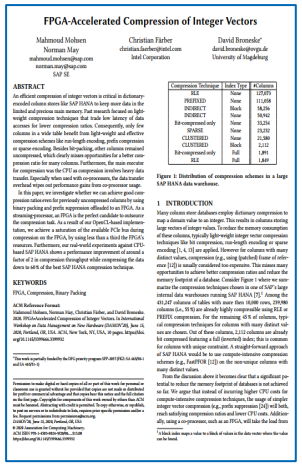
FPGA-Accelerated Compression of Integer Vectors

Mahmoud Mohsen, Norman May, Christian Färber, David Broneske

ABSTRACT

KEYWORDS

Source:
https://www.intel.de/content/www/de/de/products/details/fpga/platforms/pac/d5005.html

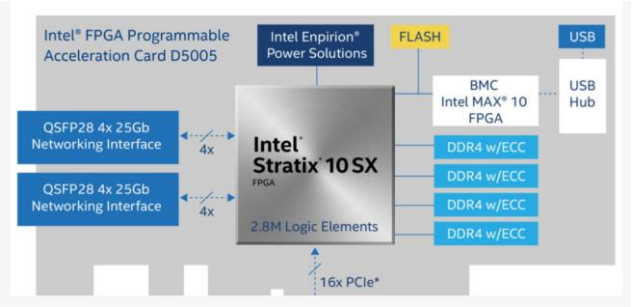

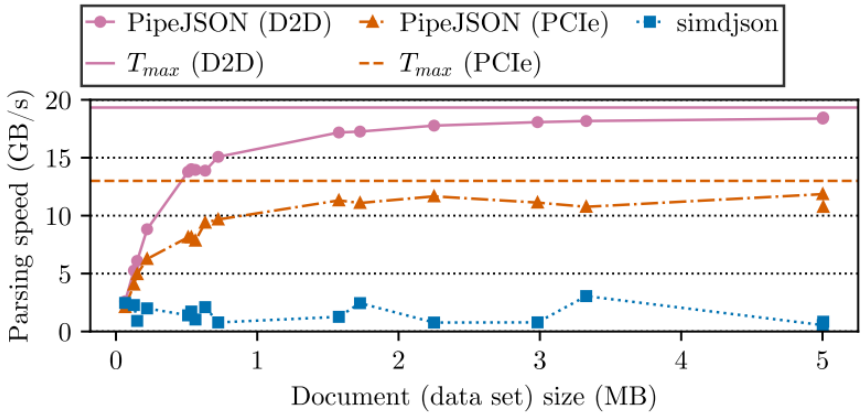

https://dl.acm.org/doi/pdf/10.1145/3533737.3535094

DAPHNE

https://www.cidrdb.org/cidr2022/papers/p4-damme.pdf

Since 2019 intel

Intel - PSG

Senior Field Application Engineer

- FAE is responsible for technical success of customer projects – customer first mentality
  - Customer trainings onside in RTL and Intel® oneAPI
  - Realizing Proof of Concepts with customers
  - Support of productization
- New feature development with Intel engineering
- Marketing support
- Find solutions for many different customers from different fields
- Working in a team of ~10 FAEs

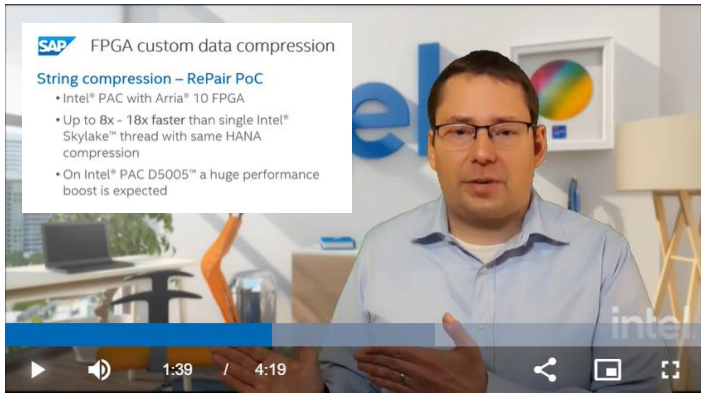

https://salesvideos.intel.com/detail/video/6234855095001/intel%C2%AE-fpga-acceleration-of-data-compression

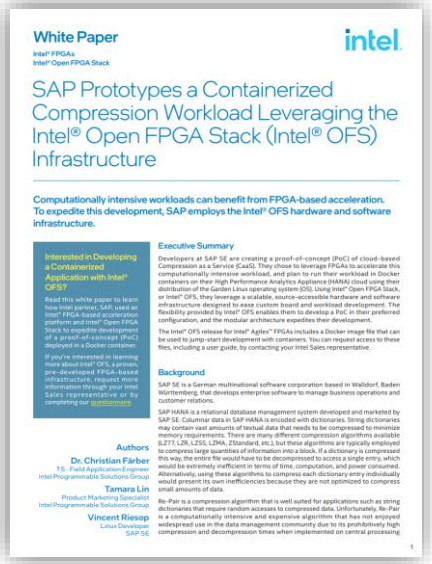
White Paper

SAP Prototypes a Containerized Compression Workload Leveraging the Intel® Open FPGA Stack (Intel® OFS) Infrastructure

https://www.intel.com.br/content/dam/www/central-libraries/us/en/documents/2022-09/sap-open-fpga-stack-white-paper.pdf

intel

Section:
What I found useful

Sub-Topics:

- Interviews
- General Thoughts
- Paths of FPGA Colleagues

Interviews I/II

What to expect

- Headhunter
  - High level, wants normally a CV and 3Qs:
    - Are you in contact with other companies already?
    - Salary (portals and talk with others)
    - When can you start?
- HR
  - Be confident
  - Rare: Some will challenge you like, why should we hire you and not the others?
  - Often: Some want to understand, what they need to do, that you join
- First time with hiring manager, job and team description, check your skills and expectations

Interviews II/II

What to expect

- Coding interview
  - Live programming, prepare: e.g. coderbyte
  - Short project, e.g. 10 days to finish
- Technical interview, often with future team members, in the past F2F
  - Prepare to talk about successful projects and small technical questions, no test
  - Be proud what you did
- Last clarification with hiring manager
- HR with last negotiation

- Ask questions!

Section:
What I found useful

Sub-Topics:

- Interviews
- General Thoughts
- Paths of FPGA Colleagues

General Thoughts

During my career path so far

- Get a mentor and mentor yourself someone
- Create a large network in and around your field
- Create an elevator pitch
- Be pro-active, take responsibility and speak up, you own your own career!
- Changing job and company is normal
- Remote work is in IT even more remote as you may think

https://www.td.org/talent-development-glossary-terms/what-is-mentoring

Section:
What I found useful

Sub-Topics:

- Interviews
- General Thoughts
- Paths of FPGA Colleagues

Paths of FPGA Colleagues

Many interesting paths

- Colleagues who did their PhD with me and worked also with FPGAs are now:
  - Group leader in engineering
    - Aerospace developing FPGA-based control systems
  - Freelancer
    - Transportation
  - ASIC designer
    - Commercial sensor chips
- Similar experience with job market and applications

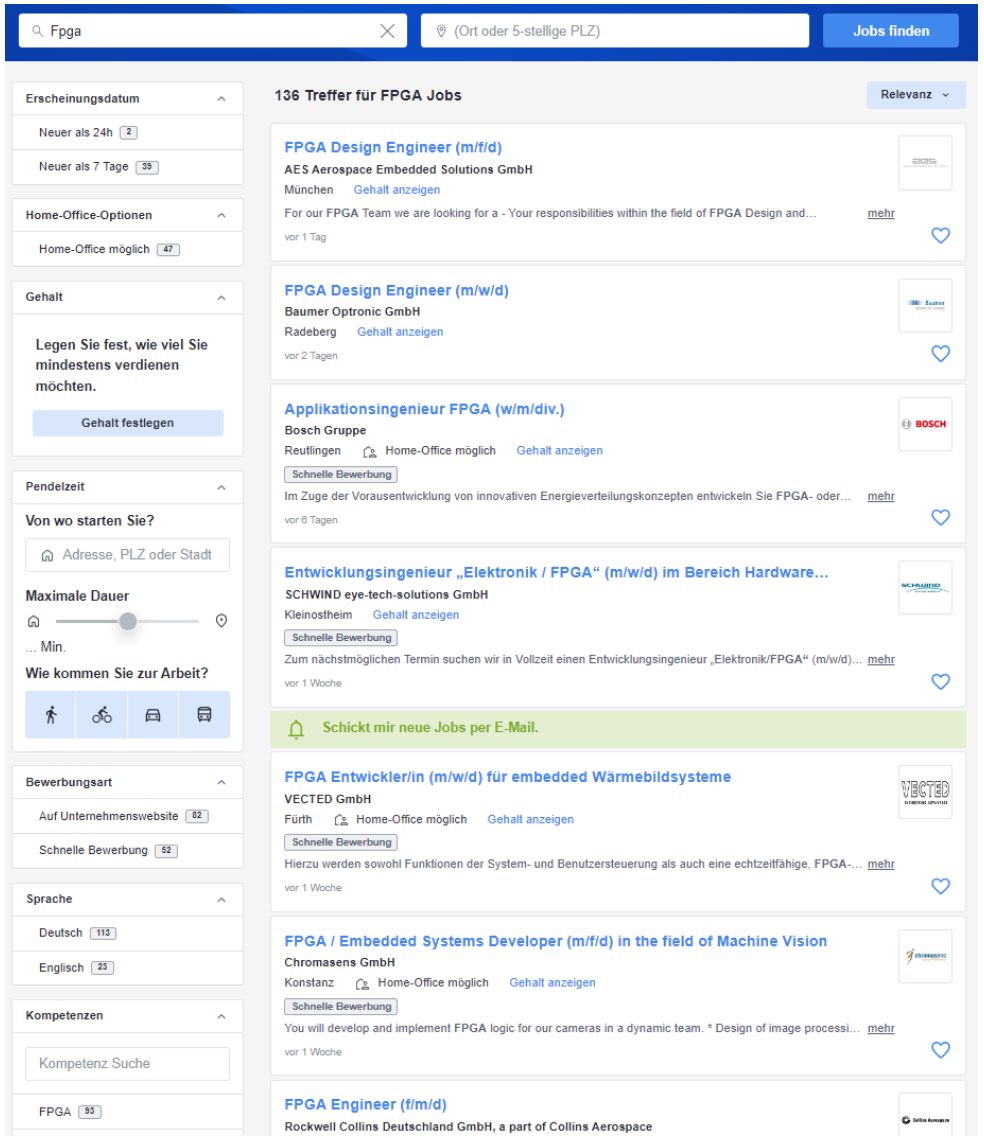

Summary

Moving from HEP to Industry with FPGAs

- Working in HEP prepared a large set of skills to be successful in industry ☺
- With hardware & FPGA knowhow you can select the job you want, great future perspective !
- FPGAs are used in a large variety of electronics from edge to cloud
  - FPGA field is growing fast
- Intel® oneAPI enables software engineers to use FPGAs without using RTL ☺
- FPGAs get used also as compute accelerators and infrastructure processing units in data centers and cloud

oneAPI

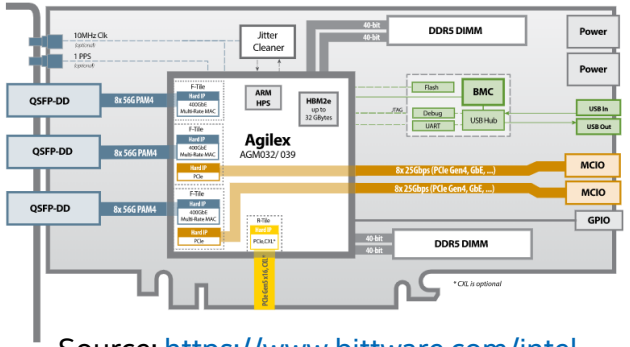

Source: https://www.bittware.com/intel

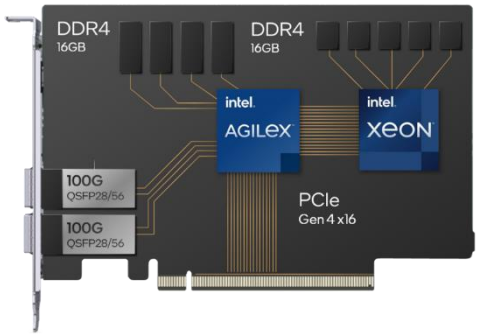

intel.

Further questions? ☺

Contact:
christian.faerber@intel.com

Or
https://www.linkedin.com/in/dr-christian-faerber-b7211590/


intel

Section:
FPGA Introduction

Sub-Topics:

- FPGA Architecture
- Use Cases

FPGA Architecture

Base Components

- ALMs (# up to 4M)
  - configurable building block that combines a look-up table and a flip-flop to implement logic functions and register-based storage.
- DSPs (# up to 12k)
  - specialized hardware block with high-speed arithmetic and signal processing capabilities used for implementing complex digital signal processing algorithms.
- RAM blocks (30MB), I/O, PLLs
- Programming: HDL (Hardware Description Language)

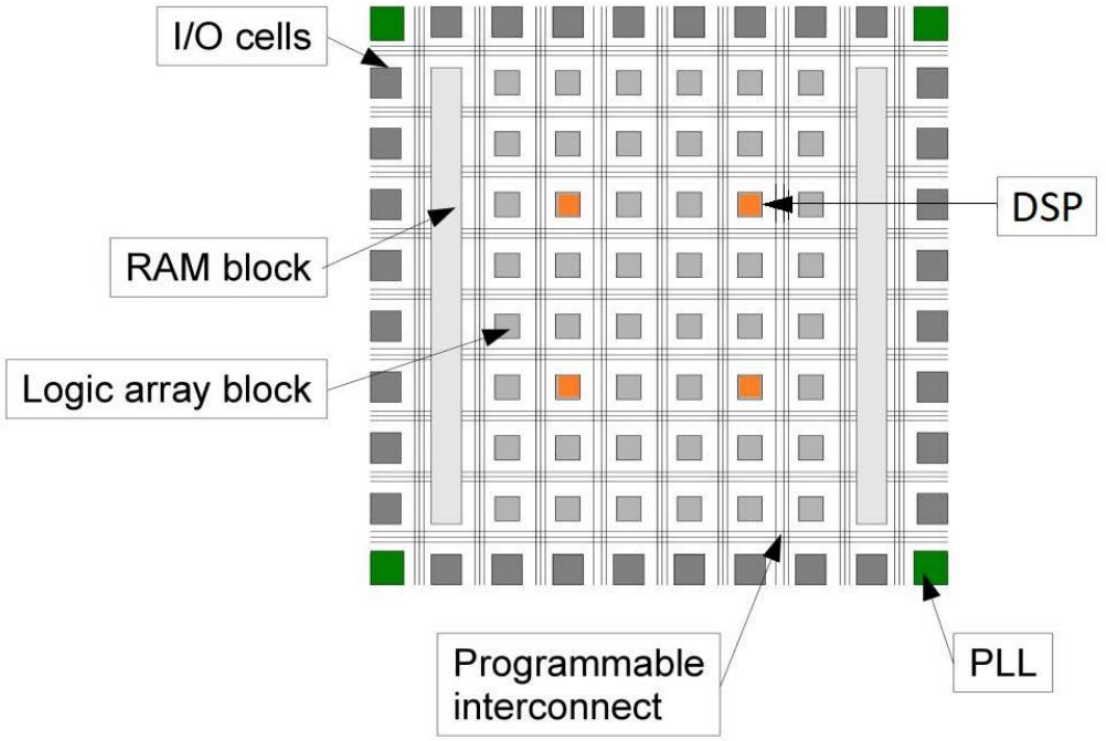

ALM

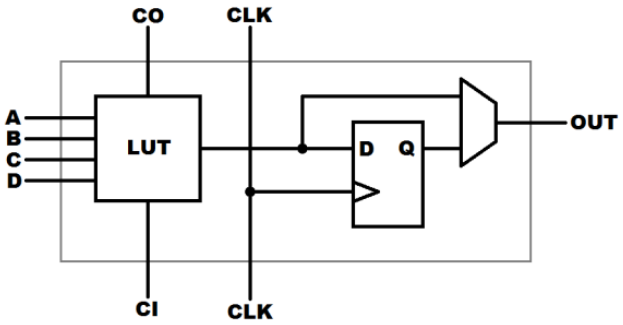

FPGA Architecture

Advanced Components

- HPS (Hard Processor System)
- Floating Point and ML DSPs
- Hyperflex registers
- Hardened IPs (e.g. DDR4, PCIe)
- High speed Transceivers
- High Bandwidth Memory
- NoC (Network on Chip)
- Chiplet tiles (PCIe5, CXL, Network, …,

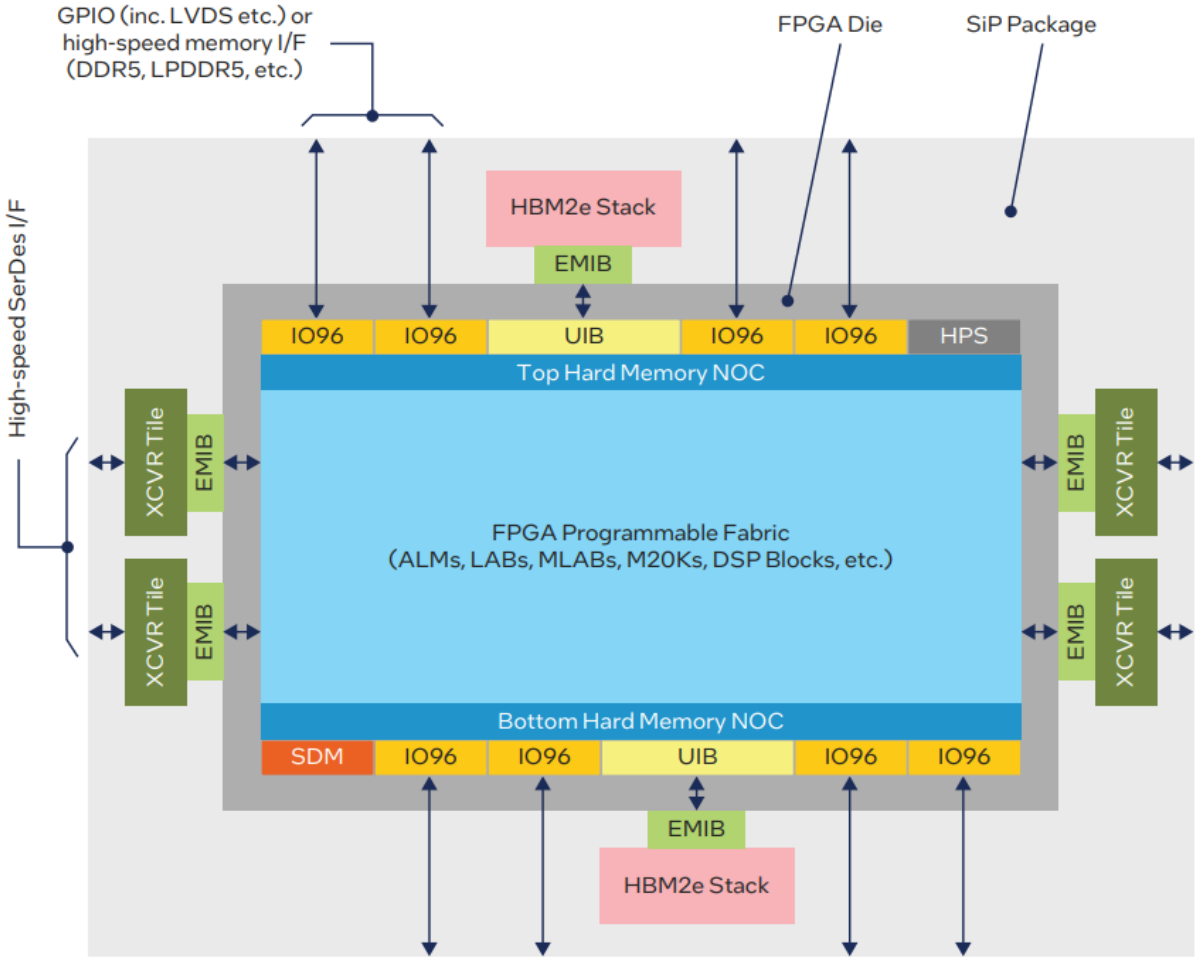

https://www.intel.de/content/www/de/de/products/docs/programmable/agilex-7-fpga-m-series-memory-bandwidth-wp.html

Section:
FPGA Introduction

Sub-Topics:

- Architecture
- Use Cases

Use Cases

FPGA base use cases

- Digital Signal Processing, Embedded Systems and Prototyping
- Aerospace, Medical, and Automotive
- Networking and Communications
- Industrial Automation and Control
- High-Performance Computing

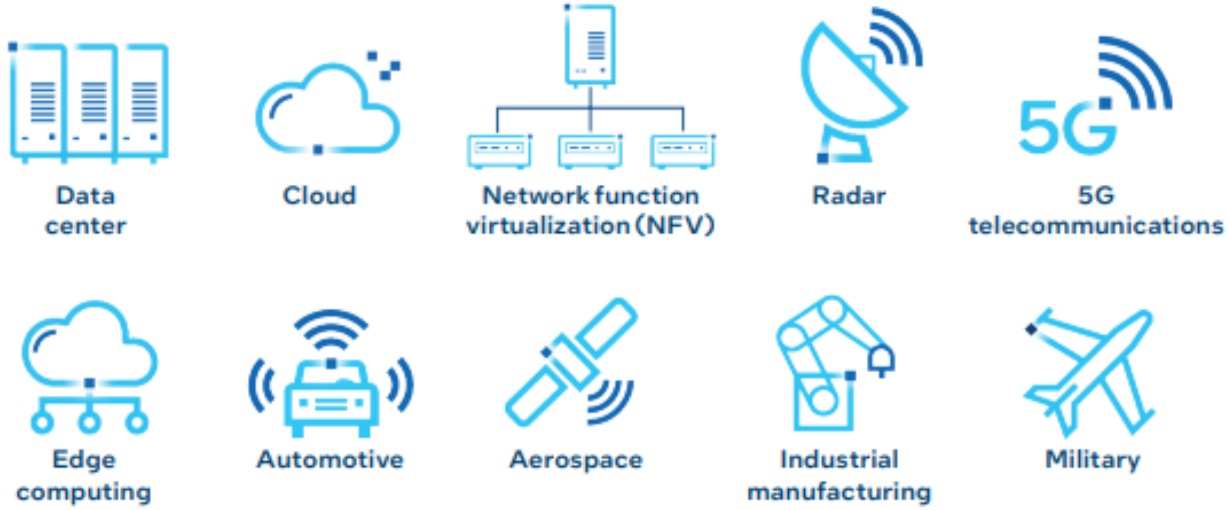

New Demands Driving Customization Needs

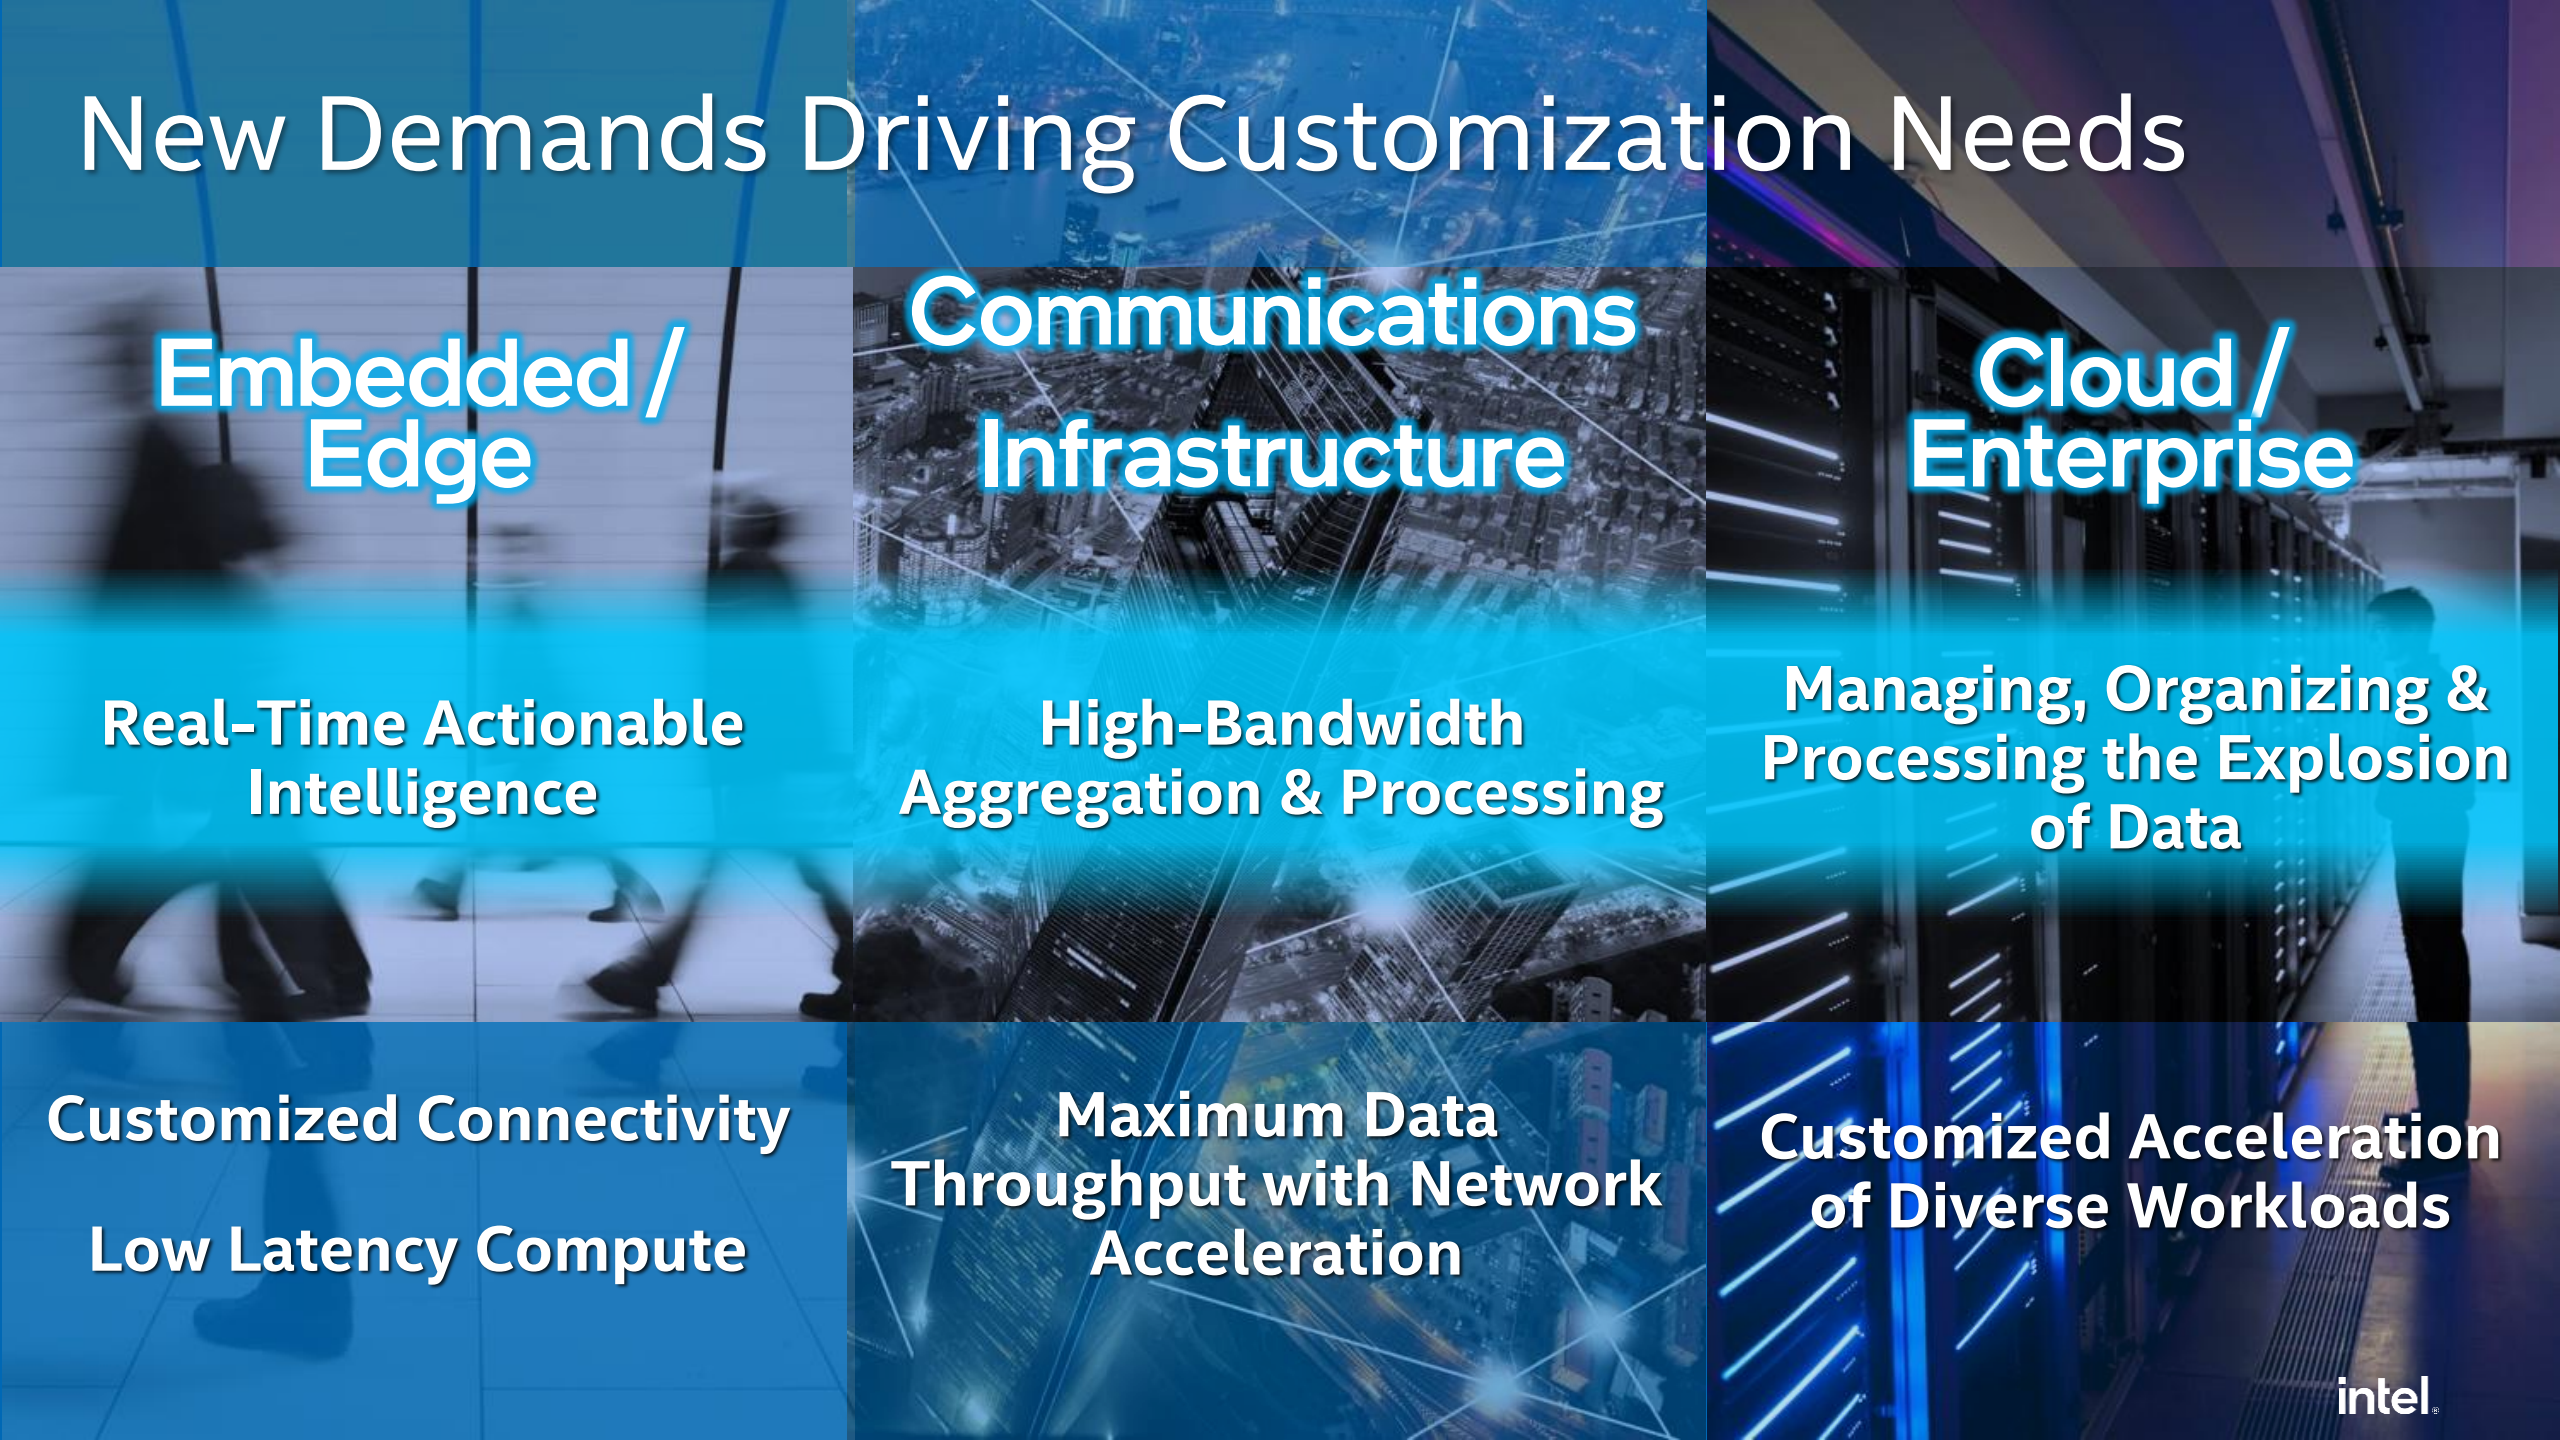

Embedded / Edge

Real-Time Actionable Intelligence

Customized Connectivity

Low Latency Compute

Communications Infrastructure

High-Bandwidth Aggregation & Processing

Maximum Data Throughput with Network Acceleration

Cloud / Enterprise

Managing, Organizing & Processing the Explosion of Data

Customized Acceleration of Diverse Workloads

intel

Section:
FPGA compute acceleration

Sub-Topics:

- Hardware
- Intel® oneAPI
- Intel® DevCloud
- FPGAs in HPC
- FPGAaaS

FPGA supported Technologies

FPGA Advantages

- UPI and now CXL
  - Cache-coherent, high bandwidth and low latency interface
  - Memory extension with DDR4/5
  - High Bandwidth Memory
  - In package 32GB HBM2e, up to 1TB/s
- IKL – FPGA to FPGA connection
- Network interfaces
- Special AI DSP blocks – S10NX
  - https://www.intel.com/content/www/us/en/products/programmable/stratix-10-nx-fpga-vs-gpu-ai-conference-paper.html

Intel® S10DX

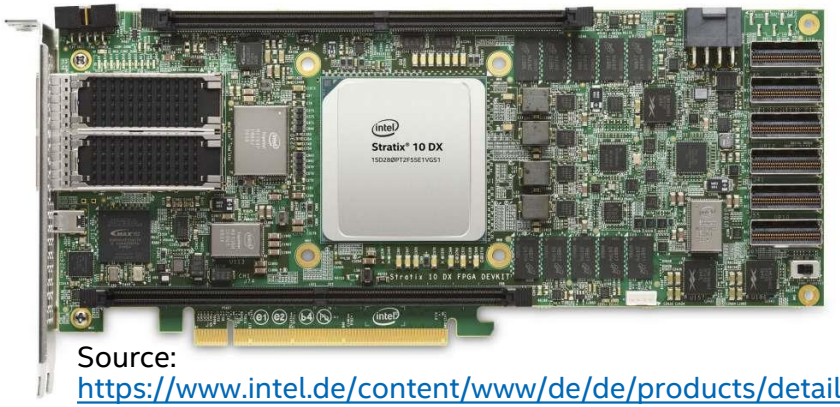

Source: https://www.intel.de/content/www/de/de/products/details/fpga/development-kits/stratix/10-dx.html

BittWare: 520NX

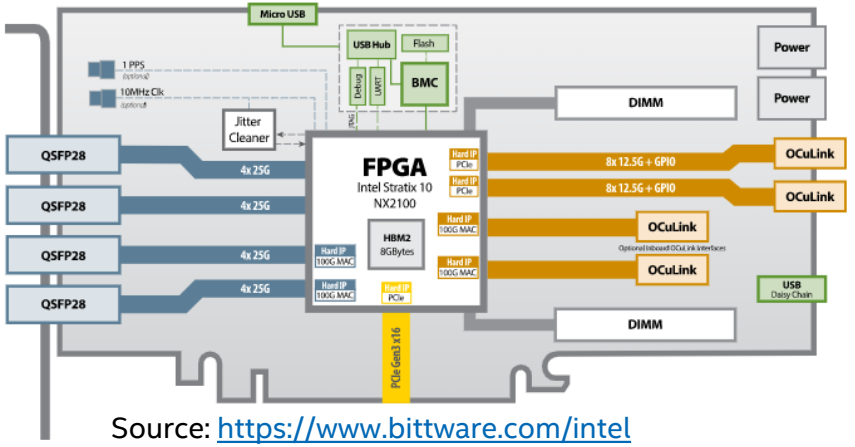

Source: https://www.bittware.com/intel

Intel® Agilex™ FPGA Cards with Intel® oneAPI

- PCIe Full Size, Dual Width Intel® Agilex™ FPGA : AGF027
- PCIe Gen4 x16
- 128GB DDR4 (4x banks)
- 3× QSFPDD (200G); 2× MCIO x8 expansion
- BittWare BMC + SDK

BittWare: IA-840f

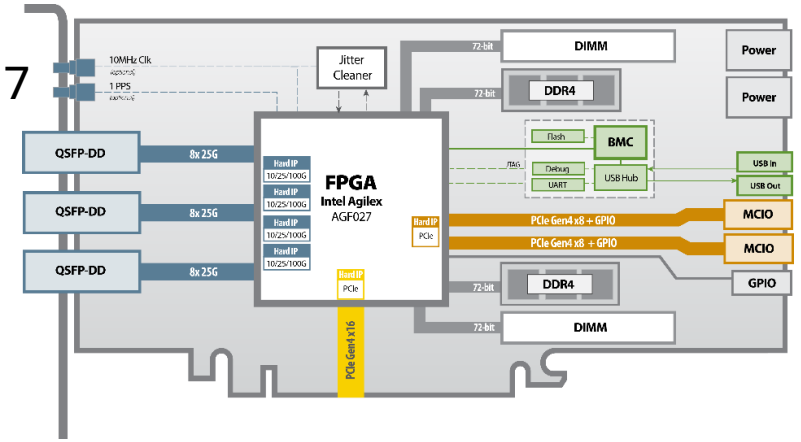

- PCIe Dual Width Intel® Agilex™ FPGA : AGM039
- PCIe Gen5 x16 + CXL
- 32GB HBM2e
- 32GB DDR5 (2x banks)
- 3× QSFP-DD (400G)
- BittWare BMC + SDK

BittWare: IA-860m

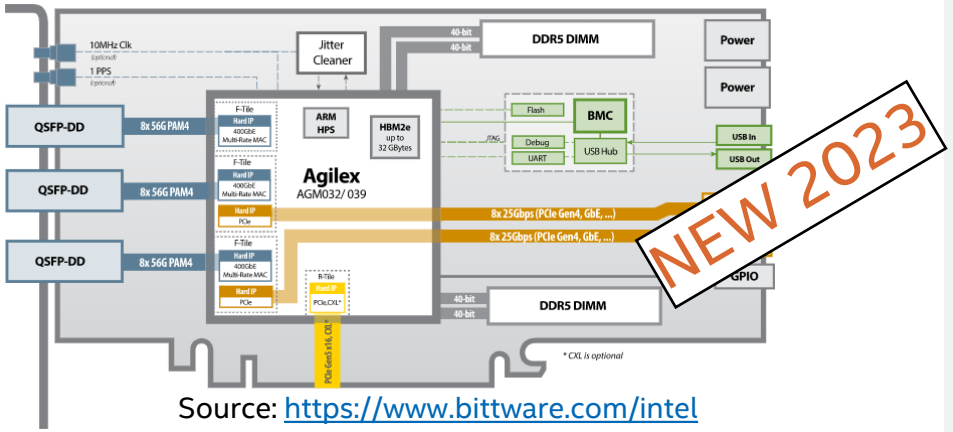

Source: https://www.bittware.com/intel

Section:
FPGA compute acceleration

Sub-Topics:

- Hardware
- Intel® oneAPI
- Intel® DevCloud
- FPGAs in HPC
- FPGAaaS

Programming Challenges

Multiple Architectures

- **Separate** programming models and toolchains for each architecture.
  - Required **training** and **licensing** – compiler, IDE, debugger, analytics/monitoring tool, deployment tool, et al. – per architecture.
  - Challenging experience in **debug**, **monitoring**, and **maintenance** of a cross-architectural source code.
  - Difficult integration across proprietary IPs and architectures and no code re-use.
- Software development complexity **limits** freedom of architectural choice.
  - Isolated investments required for technical expertise to overcome the barrier-to-entry

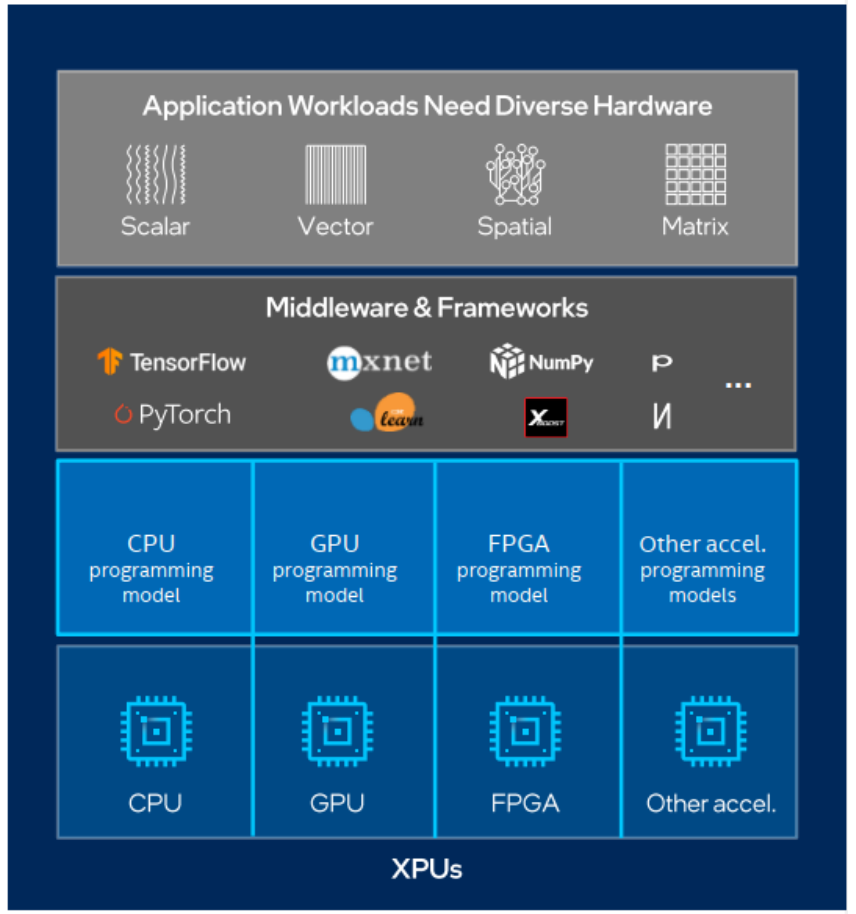

Intel® oneAPI Product

Faster Development

- Performance tuning and timing closure through **emulation** and **reports**.
- Runtime analysis via **VTune™ Profiler**
- Complex hardware patterns implemented through built-in language features: **macros**, **pragmas**, **headers**

Extensible Code

- **Code re-use** across architectures and vendors**.**
- Compatible with existing high-performance languages.

Reduced Barrier of Entry

- Leverage familiar sequential programming languages: improved **ramp-up** and **debug** time.
- IDE Integration: **Eclipse, VS, VS Code**

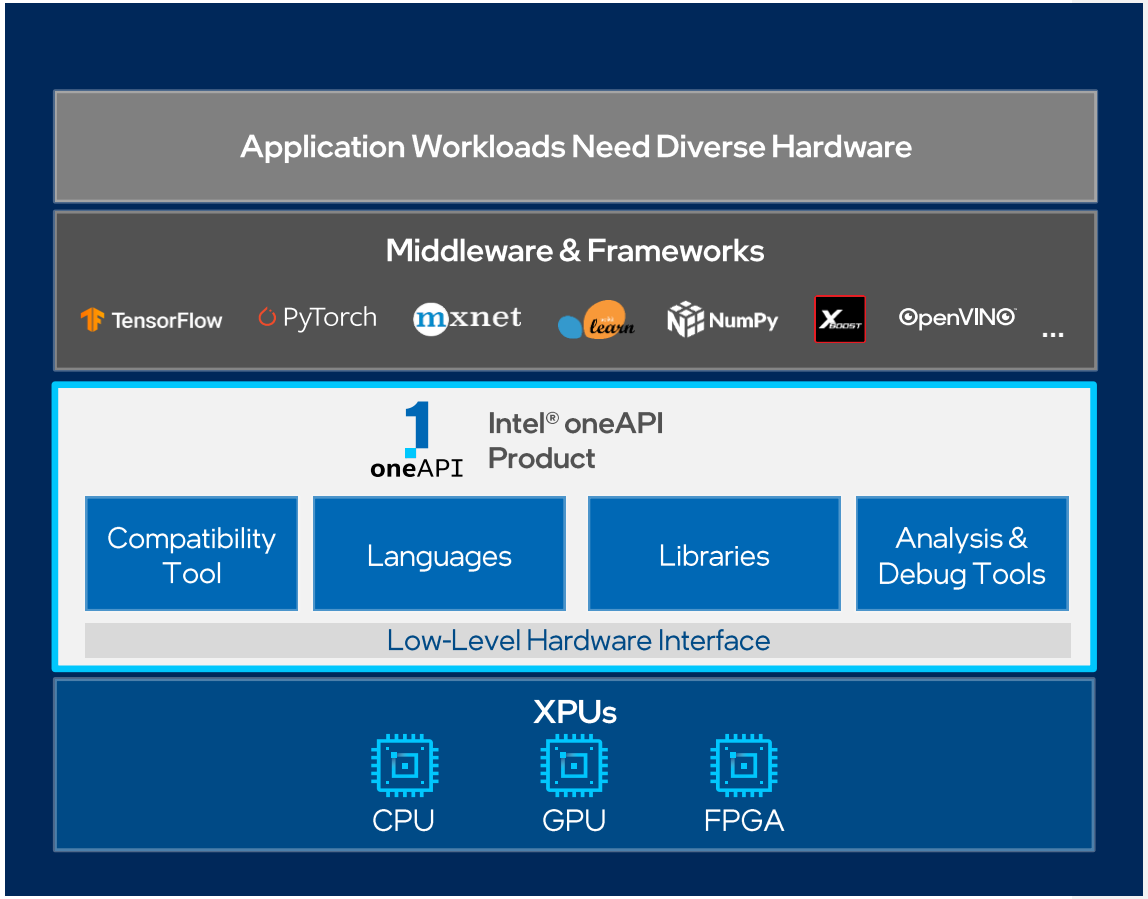

Available Now

DPC++: Three Scopes

- DPC++ Programs consist of 3 scopes:
  - Application scope - Normal host code
  - Command group scope - Submitting data and commands that are for the accelerator
  - Kernel scope – Code executed on the accelerator
- The full capabilities of C++ are available at application and command group scope
- At kernel scope there are limitations in accepted C++
  - Most important is no recursive code
  - See SYCL specification for complete list

```
void dpcpp_code(int* a, int* b, int* c) {
  //Set up an FPGA device selector
  INTEL::fpga_selector selector;
  // Set up a DPC++ device queue
  queue q(selector);
  // Setup buffers for input and output vectors
  buffer buf_a(a, range<1>(N));
  buffer buf_b(b, range<1>(N));
  buffer buf_c(c, range<1>(N));
  //Submit Command group function object to the queue
  q.submit([&](handler &h){
    //Create device accessors to buffers
    accessor a(buf_a, h, read_only);
    accessor b(buf_b, h, read_only);
    accessor c(buf_c, h, write_only);
    //Dispatch the kernel
    h.single_task<VectorAdd>([=]() {
      for (int i = 0; i < kSize; i++) {
        c[i] = a[i] + b[i];
      }
    });
  });
}
```

Application Scope

Command Group Scope

Kernel Scope

Getting Started with oneAPI on an FPGA

Intel® oneAPI Base Toolkit

+

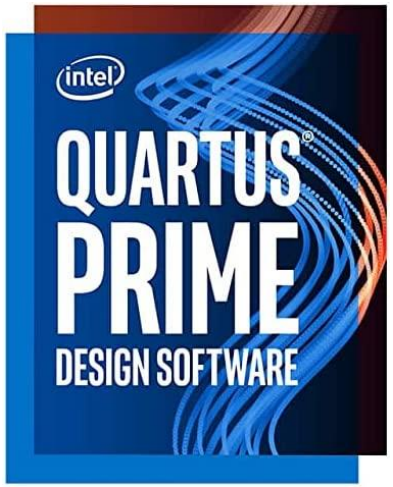

Intel Quartus® Prime Design Software

+

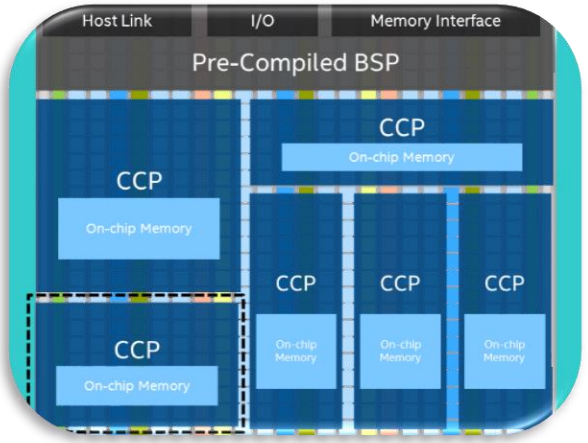

Board Support Package (BSP)

Note: Developers using custom platforms should write with IOFS their own BSP or obtain a BSP from their 3rd part platform vendor.

Section:
FPGA compute acceleration

Sub-Topics:

- Hardware
- Intel® oneAPI
- Intel® DevCloud
- FPGAs in HPC
- FPGAaaS

Intel® DevCloud

- Sign up here:
  - https://software.intel.com/devcloud
  - Account for 120 days
  - Nodes with cards installed in the group fpga_runtime
  - Nodes with extra memory for full FPGA compiles in the group fpga_compile
  - Intel® oneAPI environment already ready

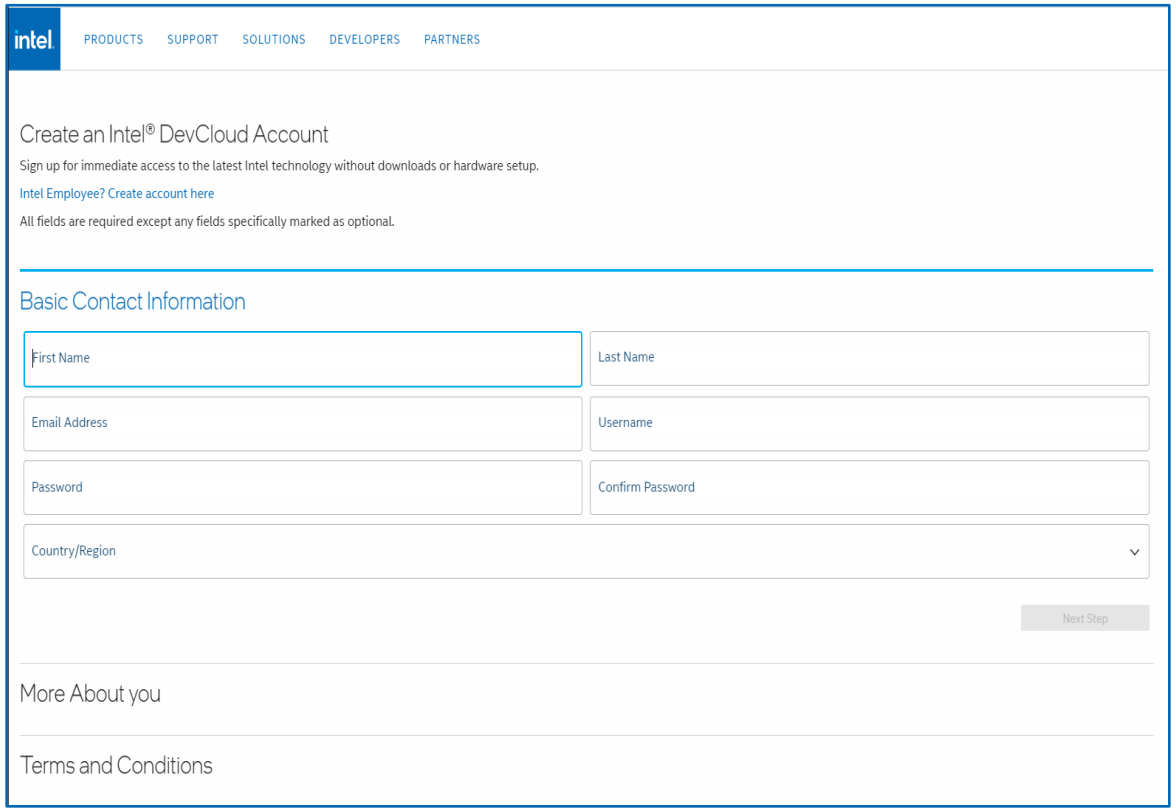

Use the Intel® DevCloud

- SSH to gateway
- Or Jupyter Notebooks via browser

- Job queue for compile and run
  - Job output into log file
- Access to single node also possible
- Session time limit:
  default : 6hrs
  max : 24hrs
- Many samples and tutorials in Git repro

https://github.com/oneapi-src/oneAPI-samples/tree/master/DirectProgramming/DPC%2B%2BFPGA

Jupyter Notebooks

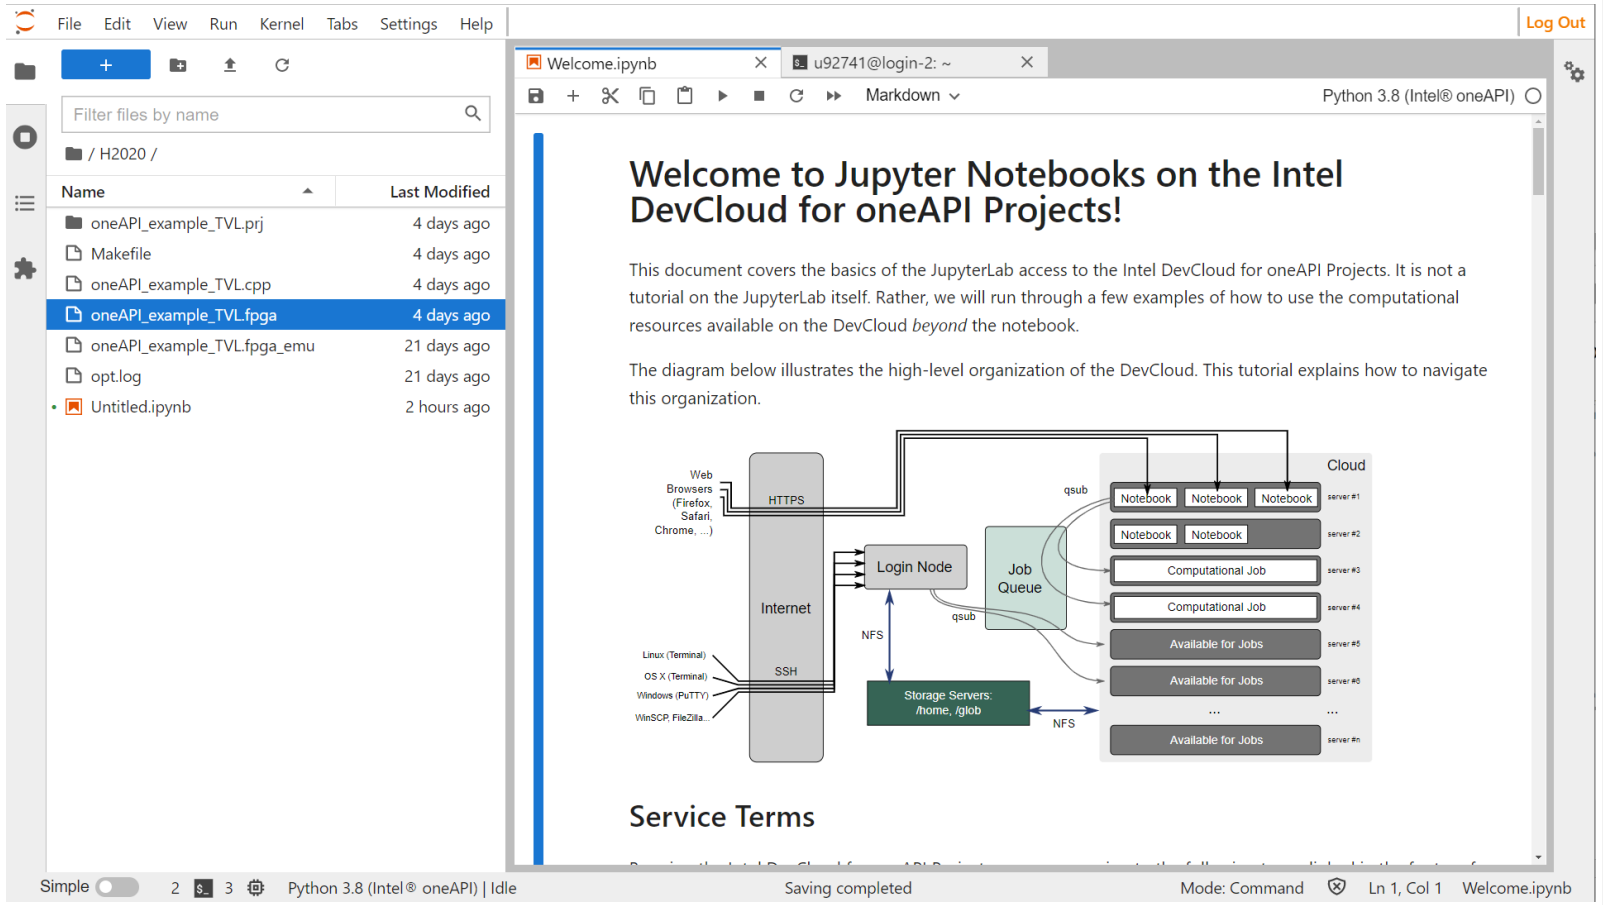

Intel® DevCloud – Available FPGA Hardware

- What are you trying to use the DevCloud for?

1) Arria™ 10 PAC - RTL AFU, OpenCL
2) Arria™ 10 - OneAPI, OpenVINO
3) Stratix™ 10 - RTL AFU, OpenCL
4) Stratix™ 10 – OneAPI
5) Emulation
6) Compilation (bitstream creation)

Soon: Intel® Agilex™ cards with Intel® oneAPI

Section:
FPGA compute acceleration

Sub-Topics:

- Hardware
- Intel® oneAPI
- Intel® DevCloud
- FPGAs in HPC
- FPGAaaS

FPGAs in HPC

Examples

- Reasons: Acceleration, Perf/W, Perf/V
- Use cases:
  - Scientific calculations
    - Genomics
    - HEP trigger
  - Data compression
  - Image processing
  - Data Base acceleration
  - File parsing
  - Financial
  - …

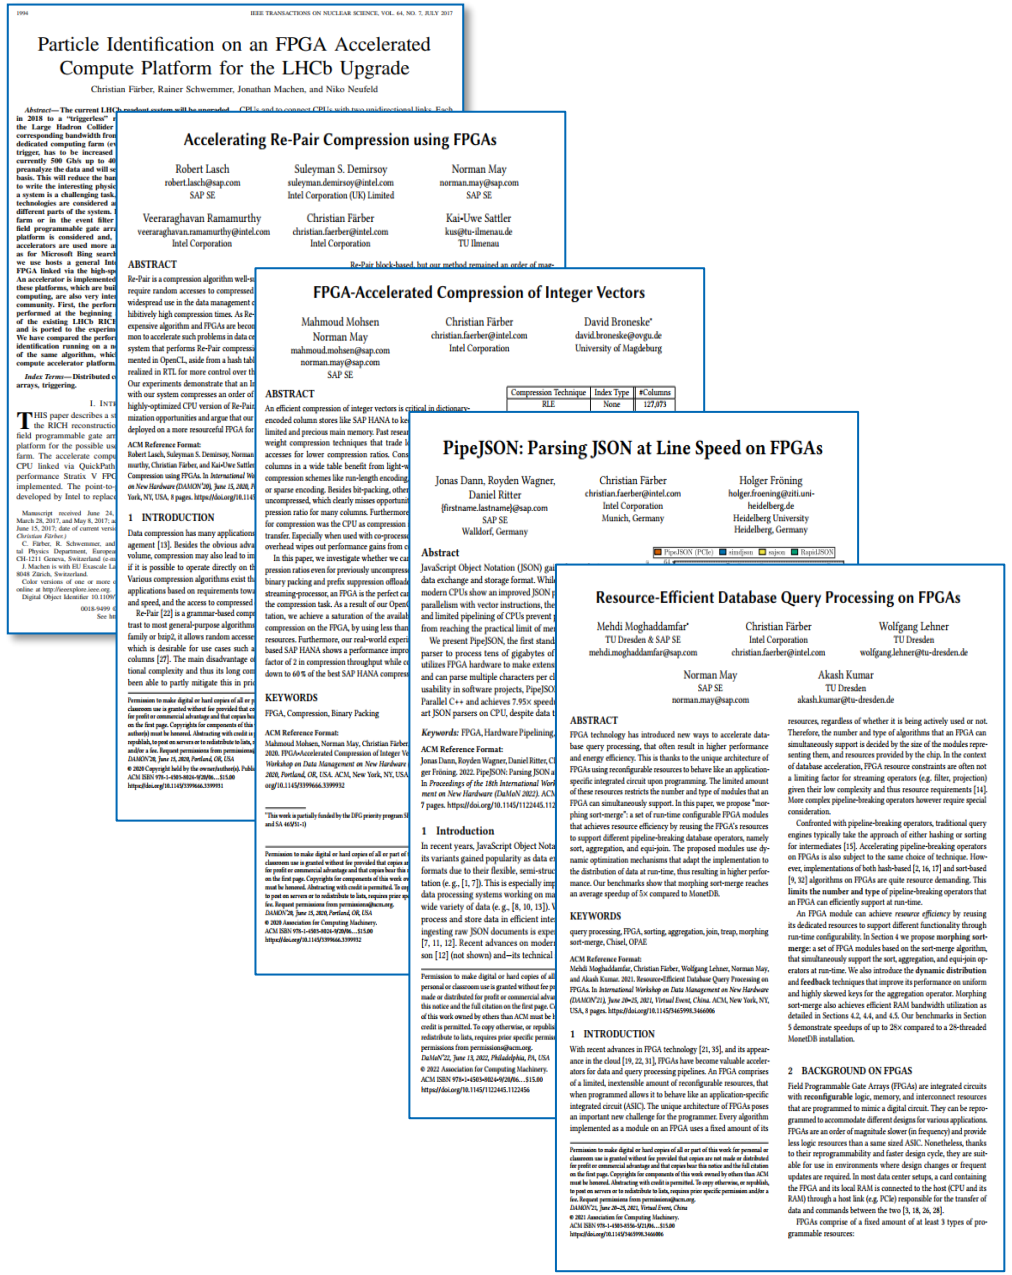

Section:
FPGA compute acceleration

Sub-Topics:

- Hardware
- Intel® oneAPI
- Intel® DevCloud
- FPGAs in HPC
- FPGAaaS

FPGAaaS

Datacenter and cloud

- FPGA accelerator functionality can be offered as a microservice enabling application developers to easily leverage many microservice characteristics:
  - Auto-deployment,
  - Scalability, dynamic configuration
- Data compression as CPU competitive example
- Sharing FPGA with multiple instances important for business case and enables even more use cases later
- IOFS enables virtualization and **Docker container** available for deployment and easy scaling, OS with DFL driver necessary e.g. **GardenLinux**
- Availability of Intel® FPGAs

CIDR 2023

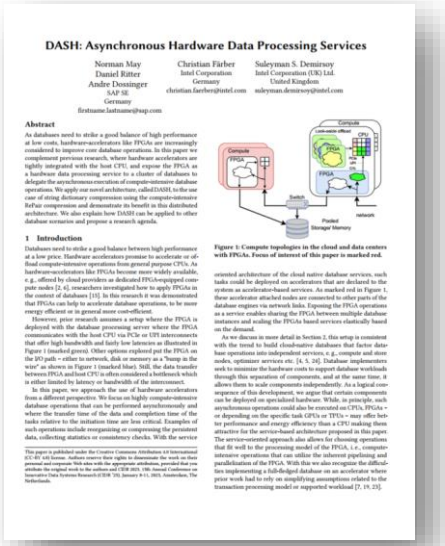

DASH: Asynchronous Hardware Data Processing Services

Norman May, Daniel Ritter, Andre Dossinger — SAP SE, Germany
Christian Färber — Intel Corporation, Germany
Suleyman S. Demirsoy — Intel Corporation (UK) Ltd., United Kingdom

Abstract

1 Introduction

https://www.cidrdb.org/cidr2023/papers/p6-may.pdf

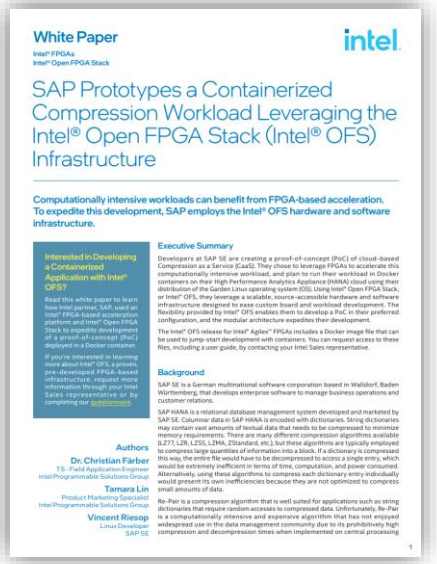

White Paper
Intel® FPGAs
Intel® Open FPGA Stack

intel

SAP Prototypes a Containerized Compression Workload Leveraging the Intel® Open FPGA Stack (Intel® OFS) Infrastructure

Computationally intensive workloads can benefit from FPGA-based acceleration. To expedite this development, SAP employs the Intel® OFS hardware and software infrastructure.

Executive Summary

Background

Authors
Dr. Christian Färber, TS - Field Application Engineer, Intel Programmable Solutions Group
Tamara Lin, Product Marketing Specialist, Intel Programmable Solutions Group
Vincent Riesop, Linux Developer, SAP SE

https://www.intel.com.br/content/dam/www/central-libraries/us/en/documents/2022-09/sap-open-fpga-stack-white-paper.pdf

Section:
Use cases

Sub-Topics:

- FPGAs and Edge Computing

FPGAs and Edge Computing

Examples

- Reduce data transfer by compute and decide at the edge
- Benefits of FPGAs:
  - Good fit for customized performant edge systems
  - Easy scalable perf., low latency, low power, thermal stability
- Applications:
  - Internet of Things (IoT)
  - Mobile and 5G
  - Industrial automation
  - Video analytics
  - Autonomous vehicles


Are FPGAs Suitable for Edge Computing?

Saman Biookaghazadeh, Arizona State University
Ming Zhao, Arizona State University
Fengbo Ren, Arizona State University

Abstract

The rapid growth of Internet-of-things (IoT) and artificial intelligence applications have called forth a new computing paradigm–edge computing. In this paper, we study the suitability of deploying FPGAs for edge computing from the perspectives of throughput sensitivity to workload size, architectural adaptiveness to algorithm characteristics, and energy efficiency. This goal is accomplished by conducting comparison experiments on an Intel Arria 10 GX1150 FPGA and an Nvidia Tesla K40m GPU. The experiment results suggest that the key advantages of adopting FPGAs for edge computing over GPUs are three-fold: 1) FPGAs can provide a consistent throughput invariant to the size of application workload, which is critical to aggregating individual service requests from various IoT sensors; (2) FPGAs offer both spatial and temporal parallelism at a fine granularity and a massive scale, which guarantees a consistently high performance for accelerating both high-concurrency and high-dependency algorithms; and (3) FPGAs feature 3–4 times lower power consumption and up to 30.7 times better energy efficiency, offering better thermal stability and lower energy cost per functionality.

1 Introduction

The Internet-of-Things (IoT) will connect 50 billion devices and is expected to generate 400 Zetta Bytes of data per year by 2020. Even considering the fast-growing size of the cloud infrastructure, the cloud is projected to fall short by two orders of magnitude to either transfer, store, or process such vast amount of streaming data [3]. Furthermore, the cloud-based solution will not be able to provide timely service for many time-sensitive IoT applications [1] [14]. Consequently, the consensus in the industry is to expand our computational infrastructure from data centers towards the edge. Over the next decade, a vast number of edge servers will be deployed to the proximity of IoT devices; a paradigm that is now referred to as fog/edge computing.

There are fundamental differences between traditional cloud and the emerging edge infrastructure. The cloud infrastructure is mainly designed for (1) fulfilling time-insensitive applications in a centralized environment; (2) serving interactive requests from end users; and (3) processing batches of static data loaded from memory/storage systems. Differently, the emerging edge infrastructure has distinct characteristics, as it keeps the promise for (1) servicing time-sensitive applications in a geographically distributed fashion; (2) mainly serving requests from IoT devices, and (3) processing streams of data from various input/output (I/O) channels. Existing IoT workloads often arrive with considerable variance in data size and require extensive computation, such as in the applications of artificial intelligence, machine learning, and natural language processing. Also, the service requests from IoT devices are usually latency-sensitive. Therefore, having a predictable performance to various workload sizes is critical for edge servers.

Existing edge servers on the market are simply a miniature version of cloud servers (cloudlet) which are primarily structured based on CPUs with tightly coupled co-processors (e.g., GPUs) [7] [6] [8] [2]. However, CPUs and GPUs are optimized towards batch processing of in-memory data and can hardly provide consistent nor predictable performance for processing streaming data coming dynamically from I/O channels. Furthermore, CPUs and GPUs are power hungry and have limited energy efficiency [4], creating enormous difficulties for deploying them in energy- or thermal-constrained application scenarios. Therefore, future edge servers call for a new general-purpose computing system stack tailored for processing streaming data from various I/O channels at low power consumption and high energy efficiency.

OpenCL-based field-programmable gate array (FPGA) computing is a promising technology for addressing the aforementioned challenges. FPGAs are

https://www.usenix.org/system/files/conference/hotedge18/hotedge18-papers-biookaghazadeh.pdf


Solution Brief
Internet of Things
intel
FPGA Cloud Connectivity Kit
Build your FPGA-based IoT application with this "cloud-ready" development kit

FPGAs can overcome the limitations of sequential processing, limited instruction sets, and basic I/O in your Internet of Things (IoT) Edge device. Intel has a complete cloud-to-edge solution based on Cyclone® V SoC FPGAs. These FPGAs make it easy for you to develop your own FPGA-based application that can collect, analyze, and react to data from cloud-connected IoT equipment.

Solution Overview

Customer Benefits

Authors
Takayuki Ikushima
Keith Woollvin

Enabled by Intel® Edge-Centric FPGAs

https://www.intel.com/content/www/us/en/products/docs/programmable/cloud-connectivity-solution-brief.html

Section:
FPGA Hardware

Sub-Topics:

- Intel® Open FPGA Stack
- CXL Interface
- FPGA based IPUs

Intel® OFS – Enabling Scale and Deployment

Intel OFS is a **software and hardware infrastructure** providing an efficient approach to develop a custom FPGA-based platform or workload using an Intel, 3rd party, or custom board.

- Scalable, source-accessible hardware and software framework delivered through **Git repositories**
- **Reduce development time** with **modular** and composable source code used as-is or easily customized
- **Upstreamed Linux kernel drivers** are being adopted by leading OS and orchestration vendors
- **Growing ecosystem** of Intel OFS-enabled boards, workloads, and OS distributions

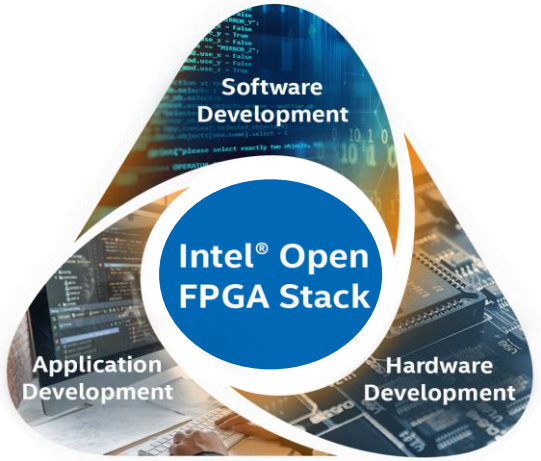

Intel® OFS Deliverables

Hardware

- Acceleration Functional Unit (AFU) Region for Workload Development with Sample AFUs
- FPGA Interface Manager (FIM)
- Board Management Controller (BMC)
- HLD enablement

Software

- Upstreamed, open-source kernel drivers
- OPAE libraries, tools and APIs
- Example Applications

Verification Environment

- UVM Verification environment provided through Git repositories

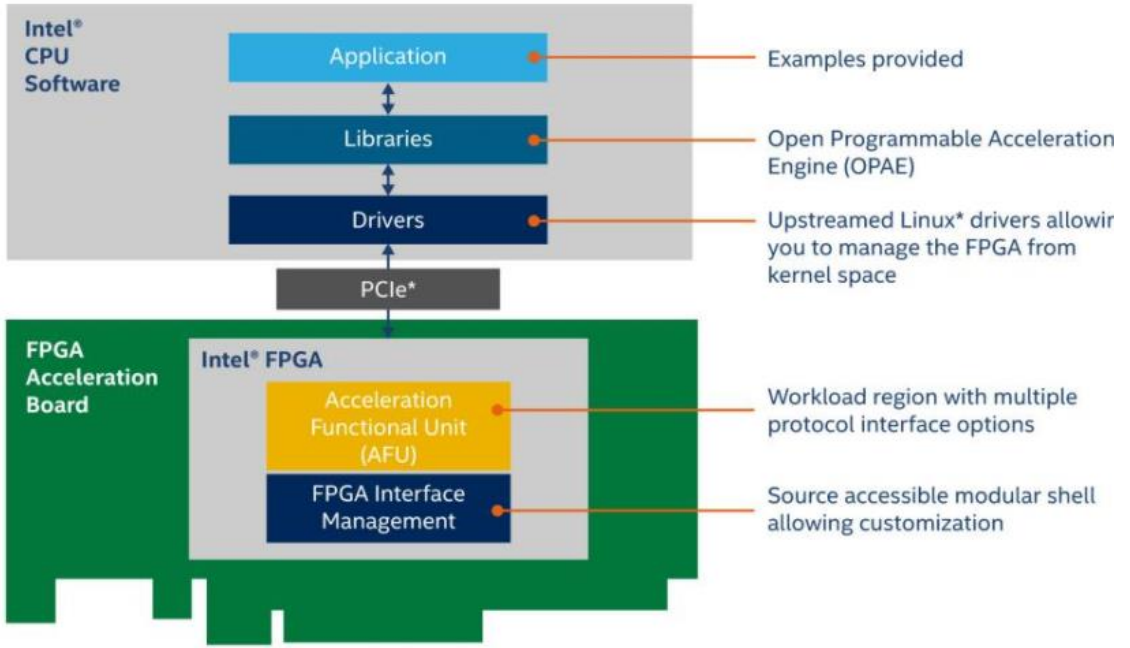

Section:
FPGA Hardware

Sub-Topics:

- Intel® Open FPGA Stack
- CXL Interface
- FPGA based IPUs

CXL Configuration Types - Type 1, 2, 3

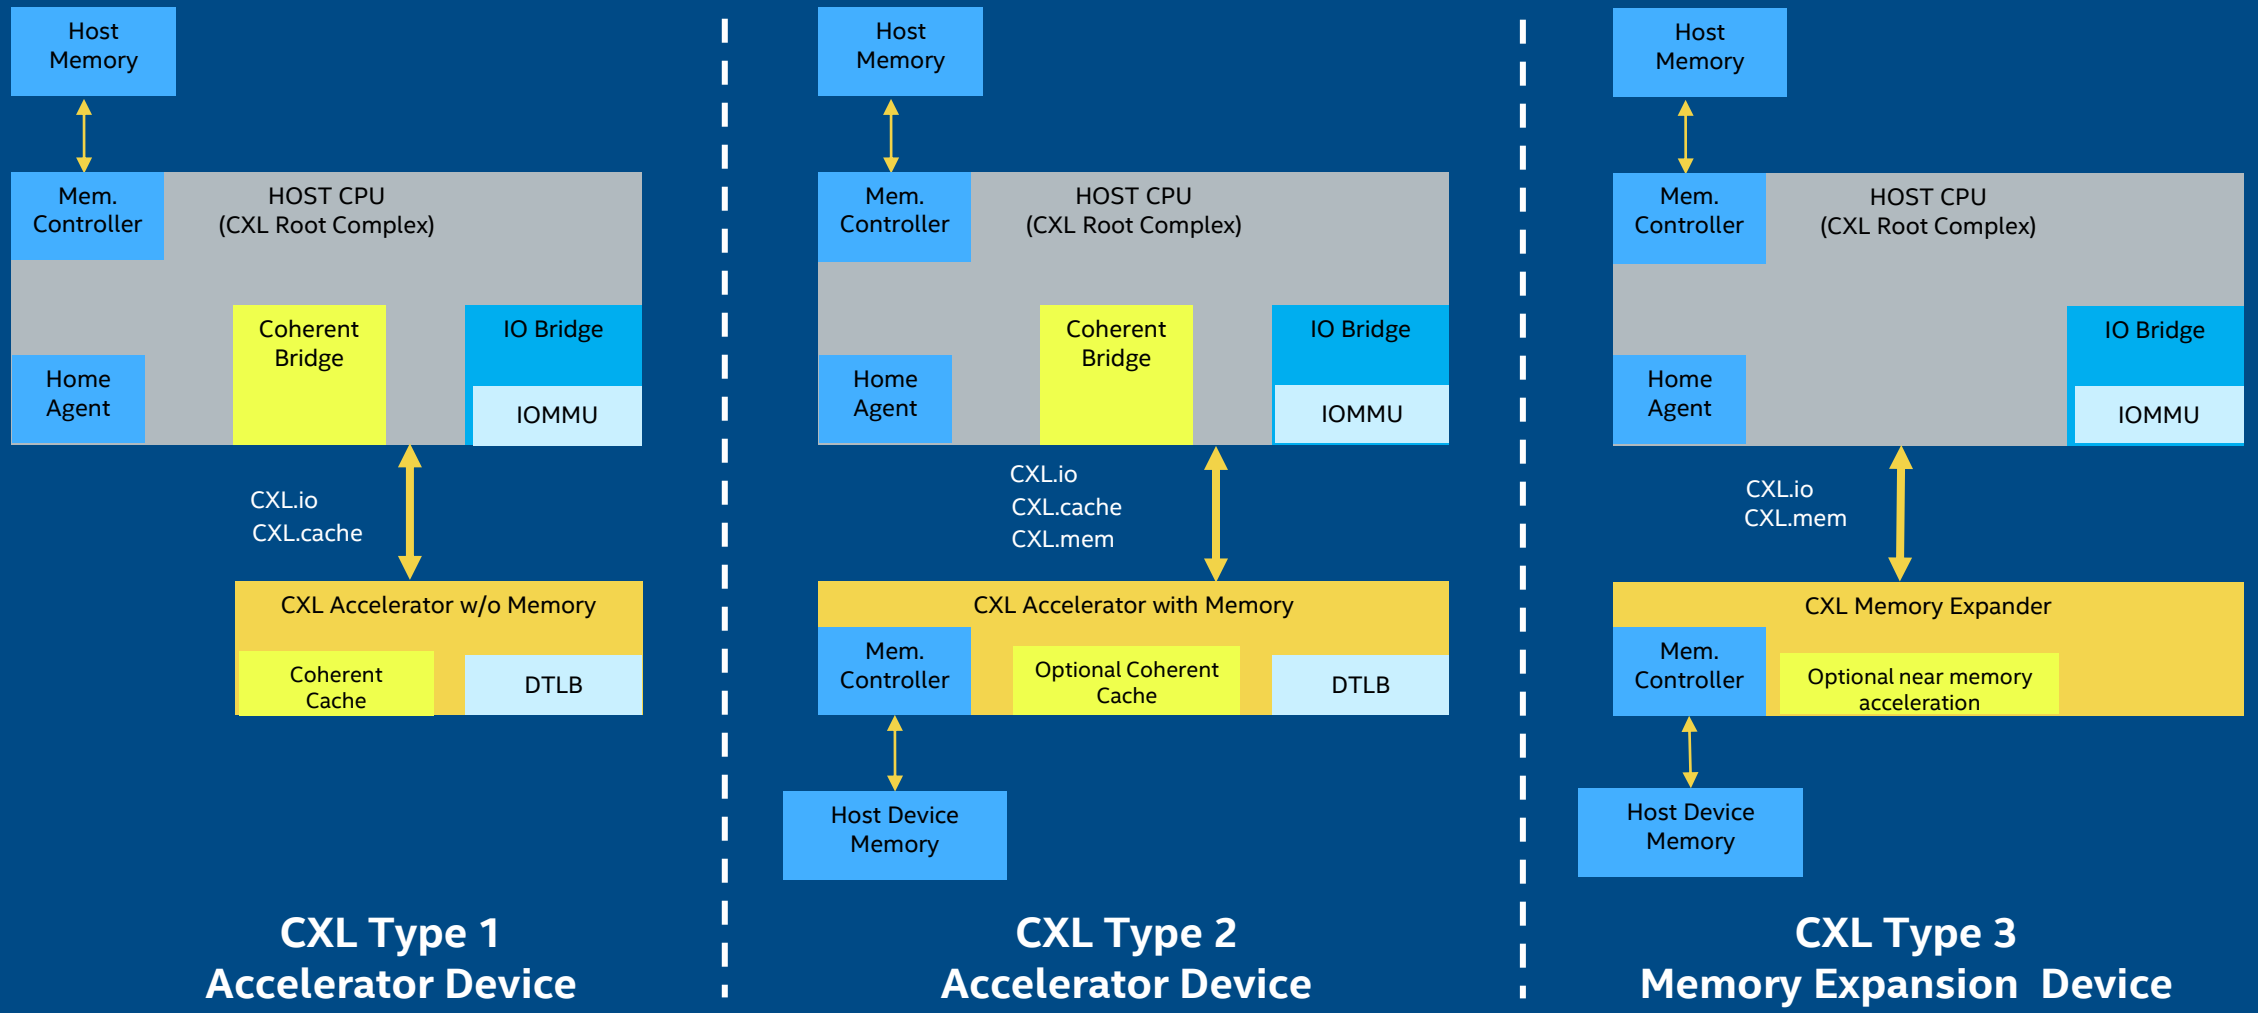

CXL - What's New?

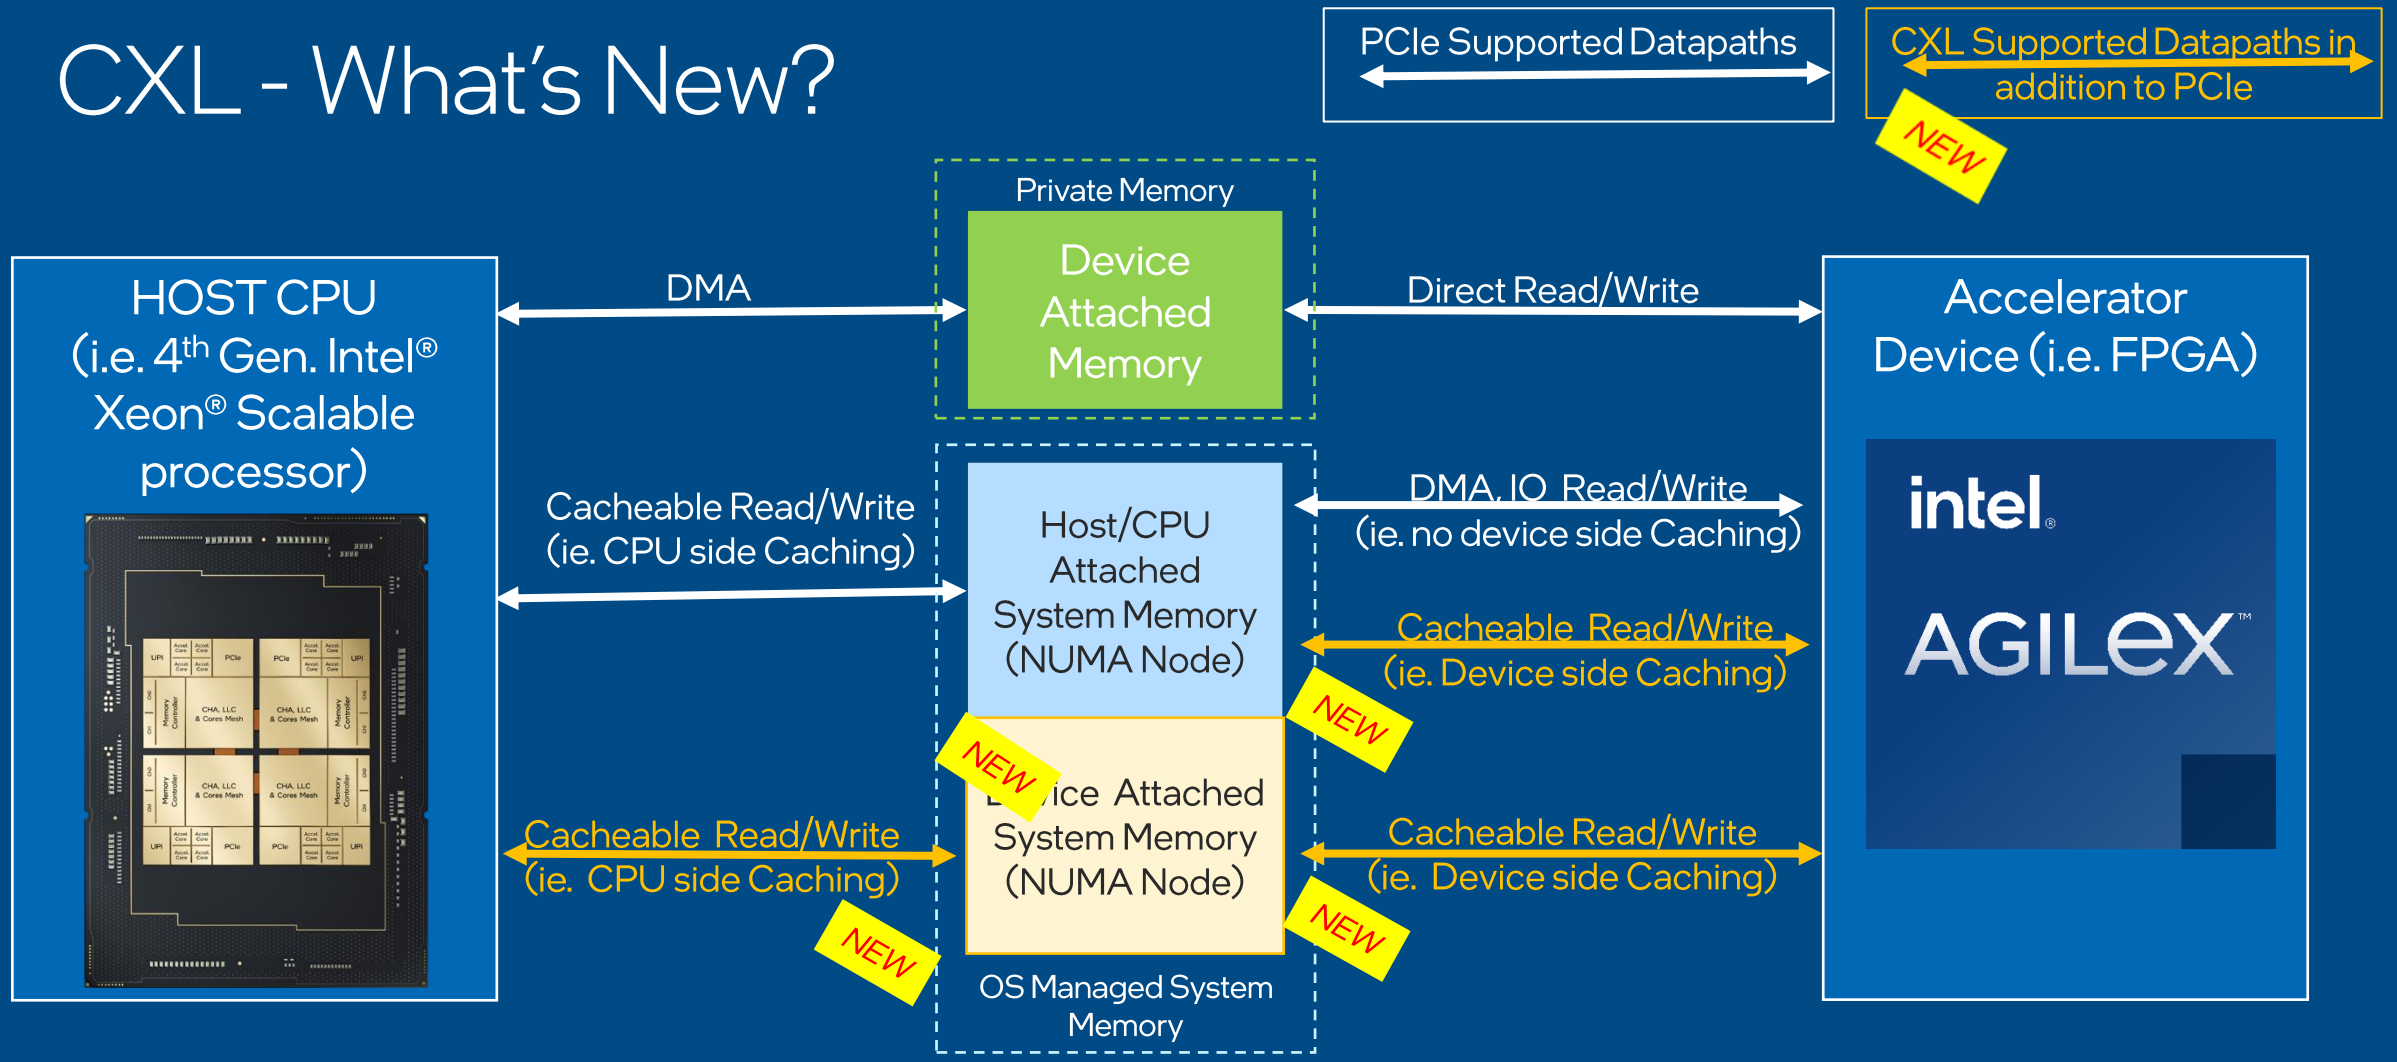

CXL-BASED ACCELERATORS – ELEVATED TO 1ST CLASS PRIORITY

Intel® Agilex™ FPGAs – The 1st FPGA with CXL Hard IP

- FPGA Industry-leading interconnect performance
  - Hard IP for CXL and PCIe 5.0 (x16 lanes)

- Supports CXL across multiple CPU / chipset suppliers
  - CXL v1.1 (now) and v2.0 (future IP release, same silicon)
  - Support for Type 1 and Type 2 Accelerators
  - Support for Type 3 memory expansion
    - Additional algorithm acceleration options may not be available from 3rd party ASSP's / ASIC's

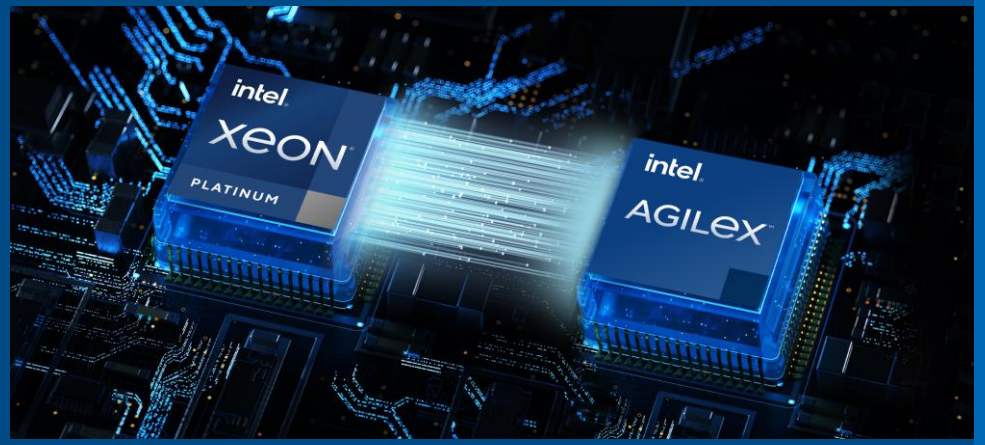

1. Intel estimates based on Intel Agilex FPGA with CXL hard/soft IP bandwidth per port vs. competitors FPGAs using 3rd party soft IP CXL controller.
2. Based on PCI-SIG integrator list results for PCIe 5.0 compliance performance. Intel Agilex R-Tile Gen 5 x16 @ 32 GT/s vs. Xilinx Versal Premium ACAP CPM5 Gen 5 x16 @ 16 GT/s.

Section:
FPGA Hardware

Sub-Topics:

- Intel® Open FPGA Stack
- CXL Interface
- FPGA based IPUs

FPGA based IPUs

Features	Target Acceleration Workloads
• 2 x 25 GbE connectivity • Intel® Stratix™ 10 DX FPGA • Intel® Xeon™ D-1612 Processor • 32GB DRAM • RTL	• Packet processing • Open vSwitch
• 2 x 100 GbE connectivity • Intel® Agilex™-F FPGA • Intel® Xeon™ D-1736 Processor • 32GB DRAM • RTL/P4	• Packet processing • Open vSwitch • NVMe-oF/RoCEV2 • RDMA

Intel® IPU Platform C5000X-PL

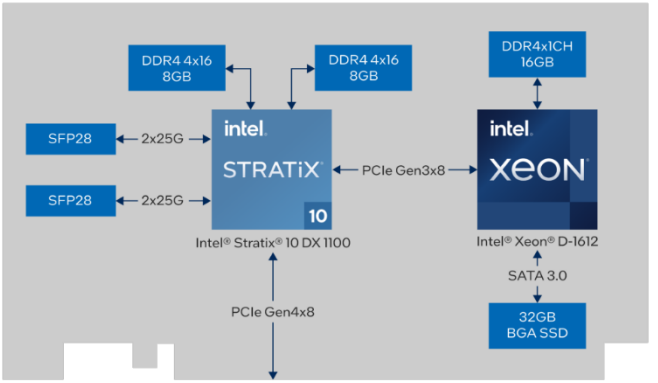

Intel® IPU Platform F2000X-PL

A coach talks to you, a mentor talks with you, and a sponsor talks about you

(roles may overlap)

COACH

A coach provides guidance for your development, often focused on soft skills (e.g., active listening) rather than technical skills.

Who Drives the Relationship?
You and your coach are responsible for driving the relationship—you can reach out to your coach when you need help, but your coach can also reach out to you.

Actions
Provide development feedback outside the formal performance evaluation process.

MENTOR

A mentor informally or formally helps you navigate your career, providing guidance for career choices and decisions.

Who Drives the Relationship?
You drive the relationship. Your mentor is reactive and responsive to your needs.

Actions
Help you determine possible career paths to meet specific career goals.

SPONSOR

A sponsor is a senior leader or other person who uses strong influence to help you obtain high-visibility assignments, promotions, or jobs.

Who Drives the Relationship?
The sponsor drives the relationship, advocating for you in many settings, including behind closed doors.

Actions
Advocate for your advancement and champion your work and potential with other senior leaders.

Based on catalyst.org guidance

@addyosmani