



AMD HACC Program Overview and Infrastructure at PC2

Christian Plessl

Paderborn University, Germany
Paderborn Center for Parallel Computing (PC2)



Paderborn
Center for
Parallel
Computing

AMD HACCC Program

Rooted in Xilinx XACC Program Announced in 2020

[Confidential - Distribution with NDA]

Program announcement 2020

[Home](#) / [News](#) / [Press Release Archive](#) / [2020 Press Releases Archive](#) / [Xilinx Teams with Leading Universities Around the World to Establish Adaptive Compute Research Clusters](#)

Xilinx Teams with Leading Universities Around the World to Establish Adaptive Compute Research Clusters

World-class research clusters at top universities to spearhead novel research into all areas of adaptive compute acceleration

May 05, 2020

World-class research clusters at top universities to spearhead novel research into all areas of adaptive compute acceleration

SAN JOSE, Calif.--(BUSINESS WIRE)-- Xilinx, Inc. (NASDAQ: XLNX), the leader in adaptive and intelligent computing, today announced it is establishing Xilinx® Adaptive Compute Clusters (XACC) at four of the world's most prestigious universities. The XACCs provide critical infrastructure and funding to support novel research in adaptive compute acceleration for high performance computing (HPC). The scope of the research is broad and encompasses systems, architecture, tools and applications.

[Confidential - Distribution with NDA]

HACCs: Heterogeneous Accelerated Compute Clusters

- Supporting High End Compute Research
- HACC community
- Remote access to Adaptive Compute hardware




Growing community of over 350 researchers at over 100 institutions

www.amd-haccs.io

[Confidential - Distribution with NDA]

Hardware resources 1.0

- Alveo Adaptive Compute hardware
 - Alveo U250 (DDR4)
 - Alveo U280 (DDR4 + HBM)
 - 100Gbps networking
- High end CPU hosts
- Vivado & Vitis software



Alveo U250




Alveo U280


[Confidential - Distribution with NDA]

Future Directions: XACC => HACC


- Heterogenous, not just FPGAs



UltraScale+ FPGA



AMD EPYC CPUs




AMD Instinct GPUs



Versal with AIE


- Scale-out



HACC Infrastructure at Paderborn Center for Parallel Computing

Supercomputer Noctua 2

- **Atos Bull Sequana XH2000**
 - 1124 servers, each with 2x AMD Milan 64-core CPUs
 - 128 NVIDIA A100 GPUs
 - **48 Xilinx Alveo U280 FPGA accelerators**
 - 32 BittWare 520N with Intel Stratix 10
 - 100 Gb/s Infiniband network
 - 6 PB DDN parallel file system
 - RedHat Enterprise Linux 8.5
- **Worldwide leading academic installation of FPGAs for HPC**
 - production system, not testbed
 - operational since April 2022



open to **HACC** users

FPGA Infrastructure in Noctua 2

	Xilinx Alveo U280 Nodes	Intel Stratix 10 Nodes	Custom Configuration Nodes
Number of Nodes	16	16	4
Accelerator Cards	3x Xilinx Alveo U280 cards	2x Bittware 520N cards	testbed for future hardware
FPGA Types	Xilinx UltraScale+ FPGA (XCU280, 3 SLRs)	Intel Stratix 10 GX 2800 FPGA	
Main Memory per Card	32 GiB DDR	32 GiB DDR	
High-Bandwidth Memory per Card	8 GiB HBM2		
Network Interfaces per Card	2x QSFP28 (100G) links	4x QSFP+ (40G) serial point-to-point links	
CPUs	2x AMD Milan 7713 , 2.0 GHz, each with 64 cores		
Main Memory	512 GiB		
Storage	480 GB local SSD in /tmp/, rest within Noctua 2 shared storage system		
Application-specific interconnect	Configurable point-to-point connections to any other FPGA connected via CALIENT S320 Optical Layer 1 Switch (OCS)		

FPGA Development and Runtime Environment on Noctua 2



- Development environment through Lmod environment modules

- Development tools `xilinx/vitis/20.2` `xilinx/vitis/21.2 (D)`
`xilinx/vitis/21.1` `xilinx/vitis_unsupported/22.1`
- FPGA shell `xilinx/u280/xdma_201920_3_2789161`
`xilinx/u280/xdma_201920_3_3246211 (D)`
- Device driver `xilinx/xrt/2.8` `xilinx/xrt/2.12 (D)`
`xilinx/xrt/2.11` `xilinx/xrt_unsupported/2.13`
- modules verify compatibility among each other, users can rely on defaults
`module load xilinx/xrt/2.12`


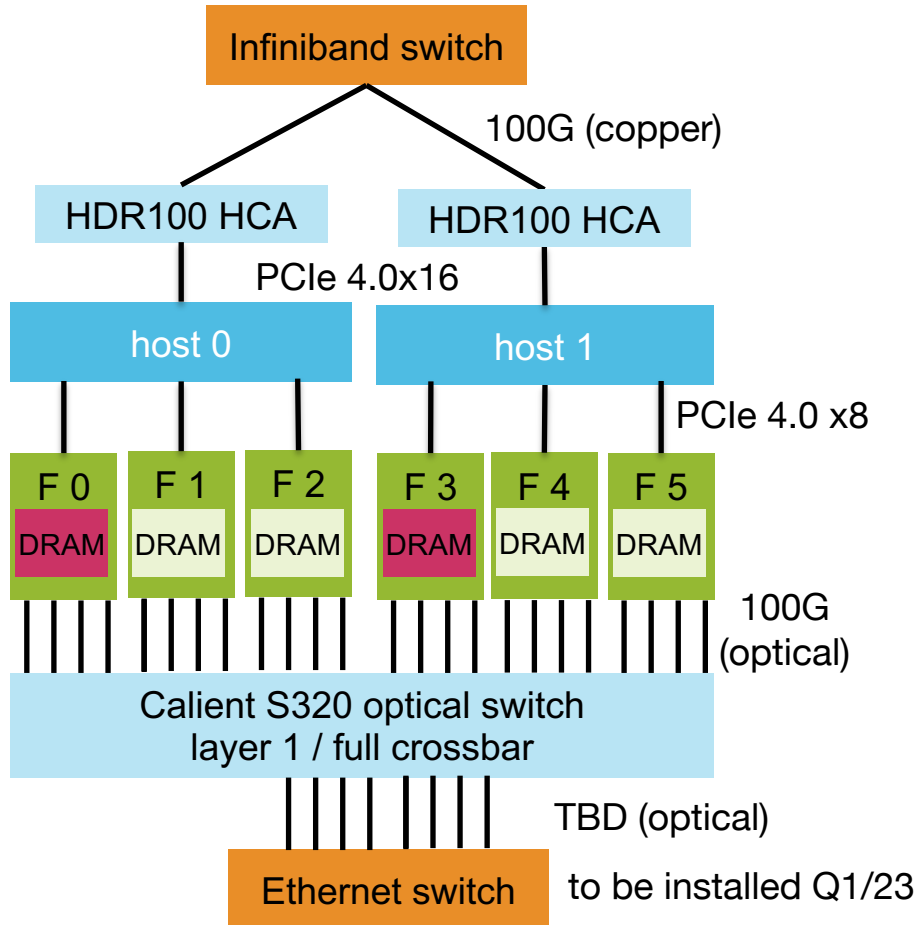
- Execution using Slurm workload manager

```
srun --partition=fpga --constraint=xilinx_u280_xrt2.12 fpga-app
```


- Users submit jobs, specify required resources
- Workload manager schedules execution on available resources



FPGA Networking: Host, FPGA-P2P L1, FPGA Switched Ethernet

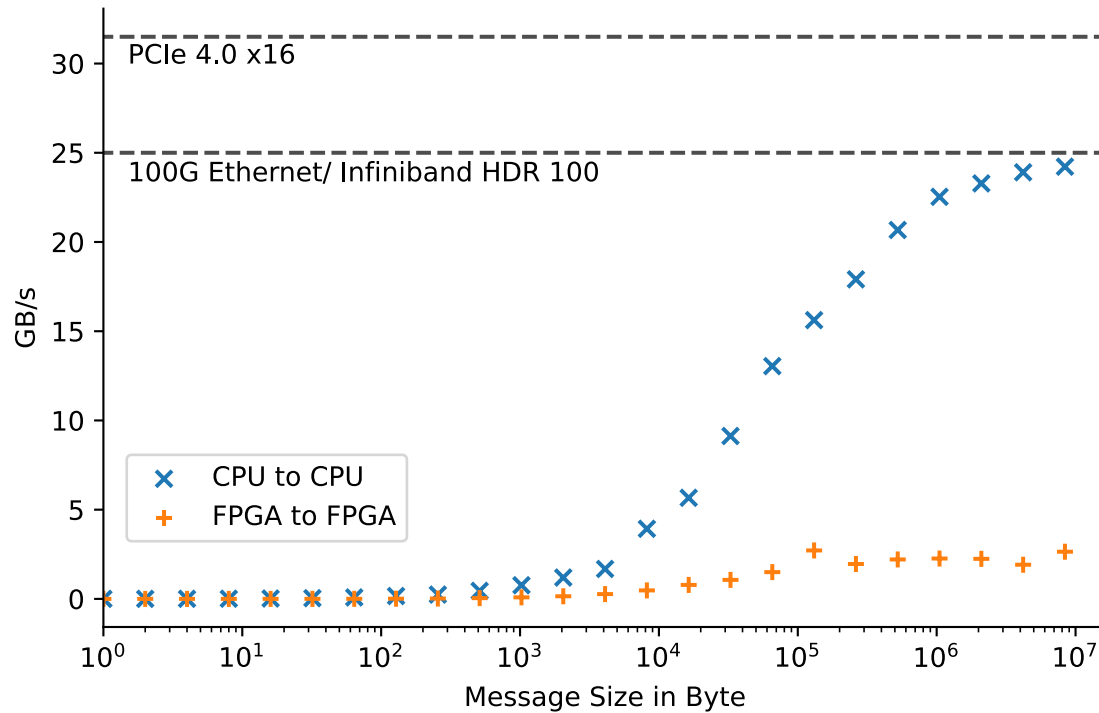


FPGA servers with optical connections

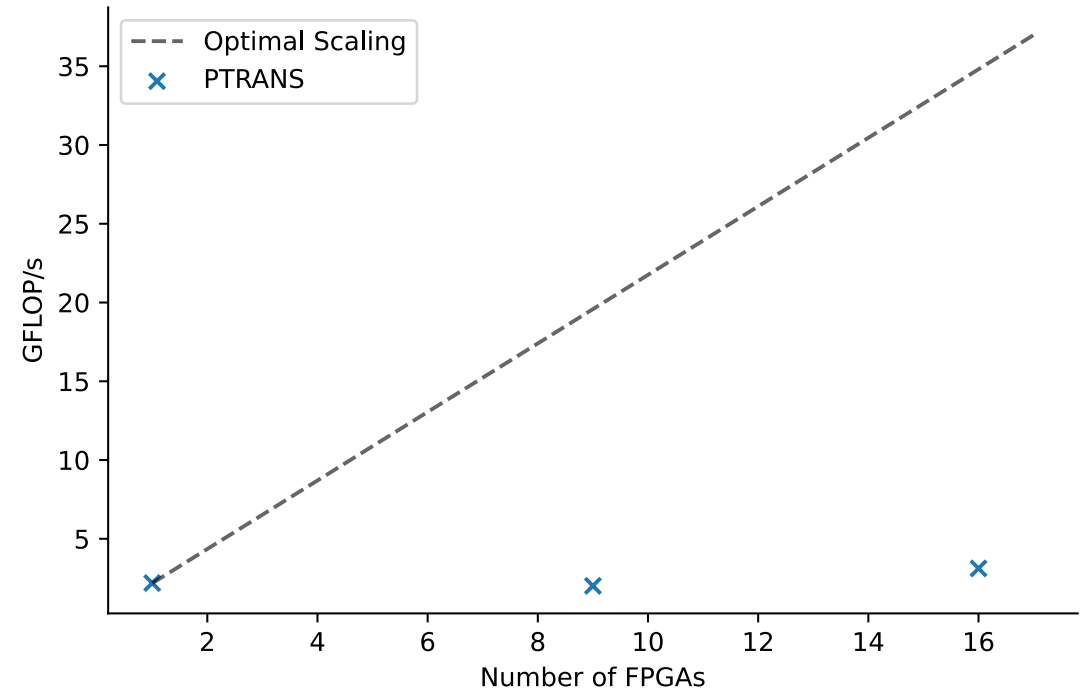


optical switch

Multi-Node Scaling Through Host – Currently Severely Limited



communication with MPI throughput via host
Infiniband (full duplex transfer)



parallel matrix transposition

How to Apply for Access in Paderborn?

- Who is eligible to access our systems
 - members of German universities (all services)
 - international users (FPGAs)
- How to apply?
 - **step 1: apply for test access**
 - quick turnaround (few days)
 - validate that infrastructure meets your needs
 - **step 2: submit formal computing time proposal**
 - use project type "small"
 - follow guides
- support by our FPGA acceleration experts and administrators with
 - application process and formalities
 - infrastructure use, optimization, best practices, ...

Small						
Projects with small resource requirements (e.g. small to medium-sized research activity, FPGA or GPU accelerated application, PhD thesis)	Noctua 2 Noctua 1	UPB, NRW (only locations without local/Tier-3 resources)	<= 4 Mio. CPU core-h and <= 40,000 GPU-h	one year	continuous submission, technical review and internal scientific review	Guide to apply for a small project
Project with a special emphasis on FPGA acceleration of HPC codes (e.g. in scope of AMD/Xilinx HACC program)	Noctua 2 Noctua 1	Global	<= 4 Mio. CPU core-h	one year	continuous submission, technical review and internal scientific review	Scientists of research and teaching institutions worldwide can be granted access to FPGA systems in the context of a research collaboration with PC ² . Guide to apply for a small project

<https://pc2.uni-paderborn.de/go/access>

Getting Started: Guides, Examples, Ready-to-use Software

- Documentation customized for available hardware and software
- “Getting started in 6 steps”
 1. Get Example Code
 2. Setup the Environment
 3. Build and Run in Emulation
 4. Create and Inspect Reports
 5. Build Hardware Design (shortcut available)
 6. Execute Design on FPGA
- References to more examples, libraries and ready-to-use software

5. Build the hardware design (bitstream)

This hardware build step (so-called hardware synthesis) can take lots of time and compute resources, so we c workload manager.

```
1 #!/bin/sh
2
3 # synthesis_script.sh
4
5 module reset
6 module load fpga
7 module load xilinx/xrt
8
9 make build TARGET=hw
```

Then, we submit the `synthesis_script.sh` to the slurm workload manager:

```
1 sbatch --partition=fpgasyn -A <your_project_acronym> -t 24:00:00 ./synthesis_script.sh
```

Details and expected output with annotations

- You can check the progress of your job via `squeue` and after the job completes, check the complete job
- Expected output

```
1 #####
2 ### Display hardware target.
3 #####
```

<https://doku.pc2.uni-paderborn.de/>

Further Information on HACCC

Further HACC Information and Application

The screenshot shows a web browser window with the URL `amd-haccs.io`. The page features a dark sidebar on the left with the AMD XILINX logo and a navigation menu with red buttons for: HACC (HACC Home, Get Started, Acknowledgment), Centers (ETH Zurich, UIUC, NUS, Paderborn, UCLA), and Resources (Events, Publications, Examples, Frameworks, Tools, Training). The main content area has a large banner with the AMD XILINX logo, the title "Heterogeneous Accelerated Compute Clusters", and the text "University Research Partners". Below this are logos for ETH zürich, ILLINOIS, NUS (National University of Singapore), PADERBORN UNIVERSITY, and UCLA. The text below the banner describes the HACC program as a special initiative for adaptive compute acceleration in HPC, mentioning its previous name XACC. It states that five HACCs are established at prestigious universities and are equipped with the latest Xilinx hardware and software. A link to the "getting started" page is provided. At the bottom, there is a "Contribute to this site" section with a "Contribute" button and a paragraph explaining how to contribute via GitHub or by contacting the Xilinx University Program team.

www.amd-haccs.io

Contact and Further Information



Contact

Christian Plessl

christian.plessl@uni-paderborn.de

Twitter: @plessl // @pc2_upb

Further information

<https://www.amd-haccs.io>

<https://doku.pc2.uni-paderborn.de>

<https://pc2.uni-paderborn.de>



Paderborn
Center for
Parallel
Computing