FBench – Defining and Optimizing OpenCL Benchmarks for FPGAs

Tobias Kenter, Christian Plessl, Marius Meyer

High-Performance IT Systems group

Paderborn University, Germany
Paderborn Center for Parallel Computing

Project Group Presentation, Paderborn, 8 July 2019

Paderborn Center for Parallel Computing

Benchmarks (1)

- Characterize system performance
- Make informed purchase decisions

Know your workload?

➢ Look up the numbers!

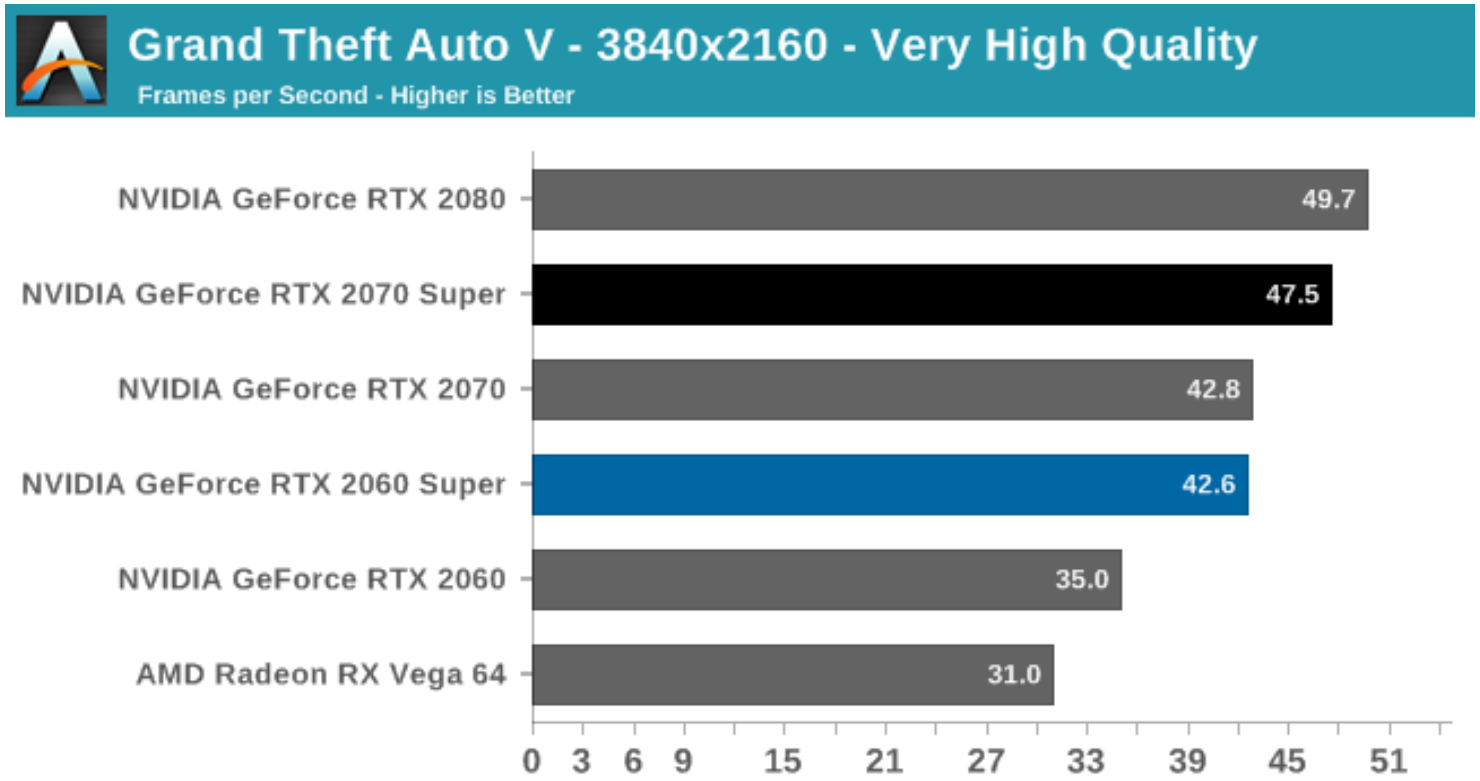

https://www.anandtech.com/show/14586/geforce-rtx-2070-super-rtx-2060-super-review/11

Benchmarks (2)

Don’t know your workload?

- Checkout all results
- Weight them in some way…

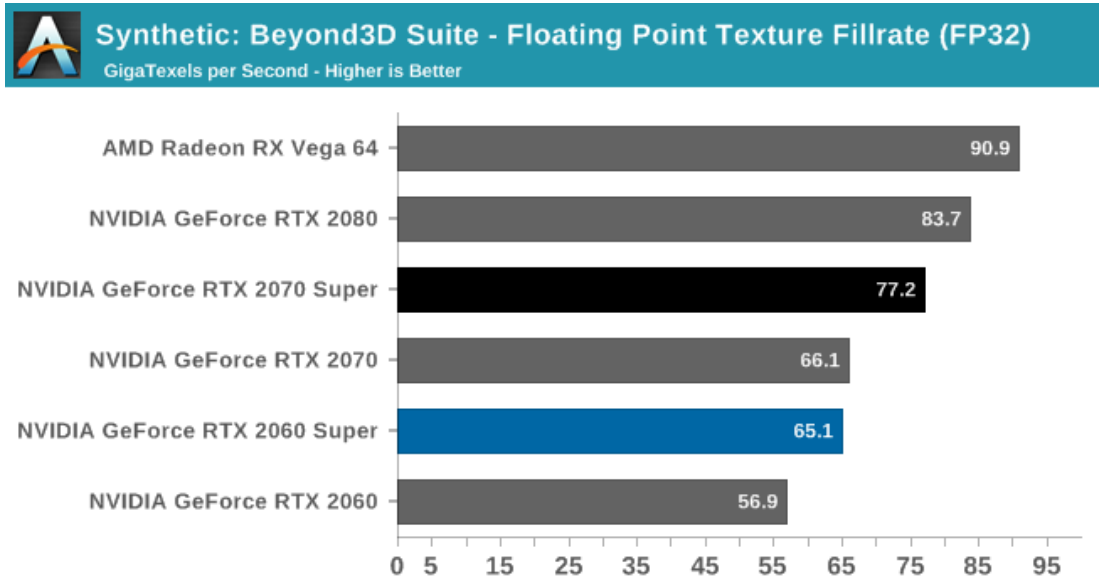

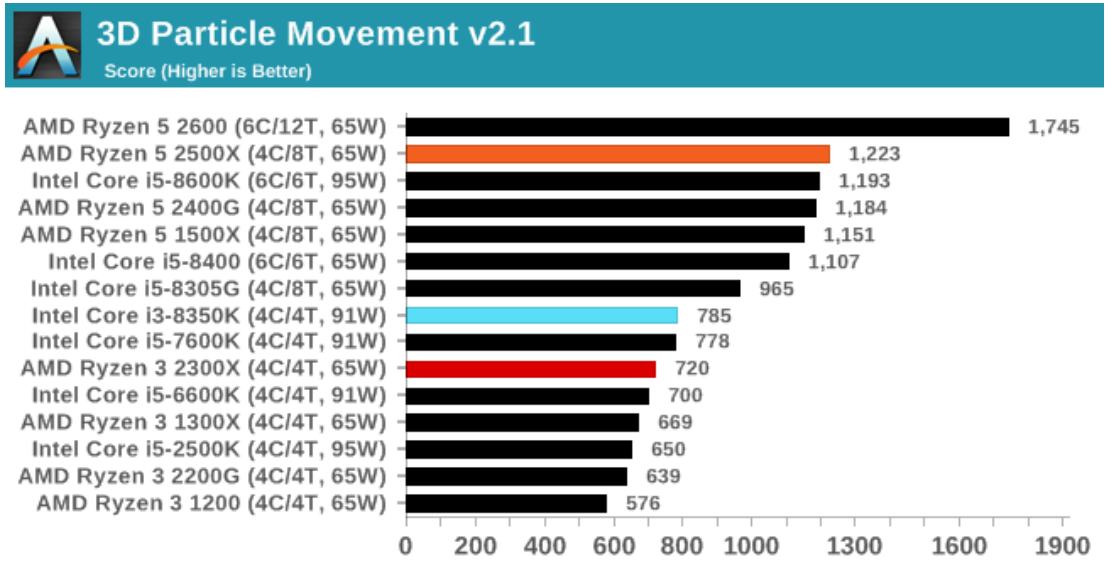

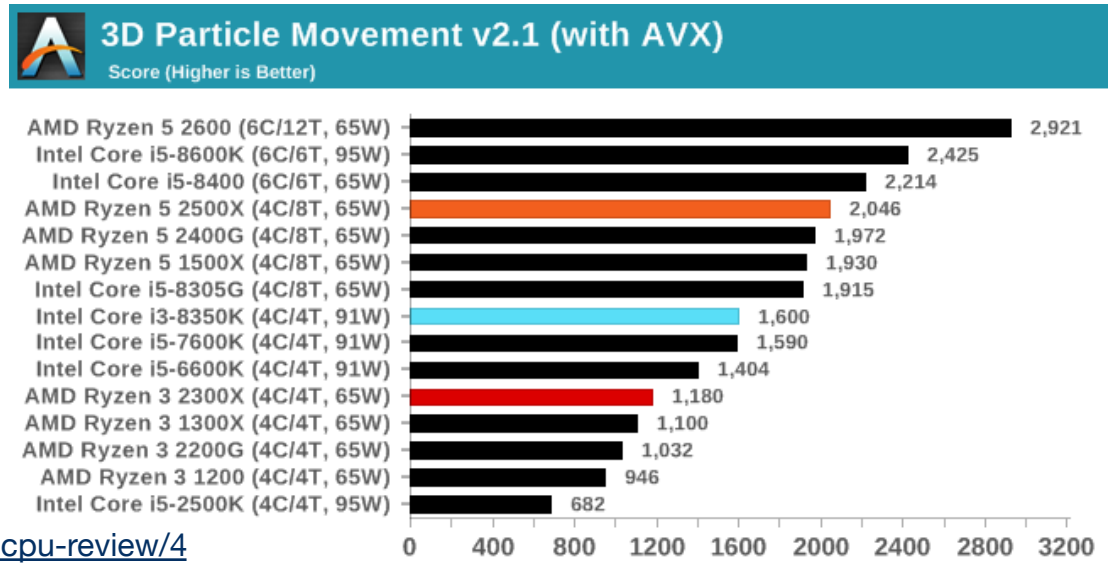

https://www.anandtech.com/show/13945/the-amd-ryzen-5-2500x-and-ryzen-3-2300x-cpu-review/4
https://www.anandtech.com/show/14586/geforce-rtx-2070-super-rtx-2060-super-review/14

Benchmarks in HPC (1)

Standard benchmarks for Top500 Lists

- High Performance Linpack (HPL)
- High-Performance Conjugate Gradient (HPCG)

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
3	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 , NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
5	**Frontera** - Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR , Dell EMC Texas Advanced Computing Center/Univ. of Texas United States	448,448	23,516.4	38,745.9	

HPCG List for June 2019

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	HPCG (TFlop/s)
1	1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	2925.75
2	2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	1795.67
3	20	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect , Fujitsu RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	602.74
4	7	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , Cray Inc. DOE/NNSA/LANL/SNL United States	979,072	20,158.7	546.12
5	8	**AI Bridging Cloud Infrastructure (ABCI)** - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR , Fujitsu National Institute of Advanced Industrial Science and Technology (AIST) Japan	391,680	19,880.0	508.85

https://www.top500.org/lists/2019/06/
https://www.top500.org/hpcg/lists/2019/06/

Benchmarks and System Characteristics

Standard benchmarks don't cover all important system characteristics

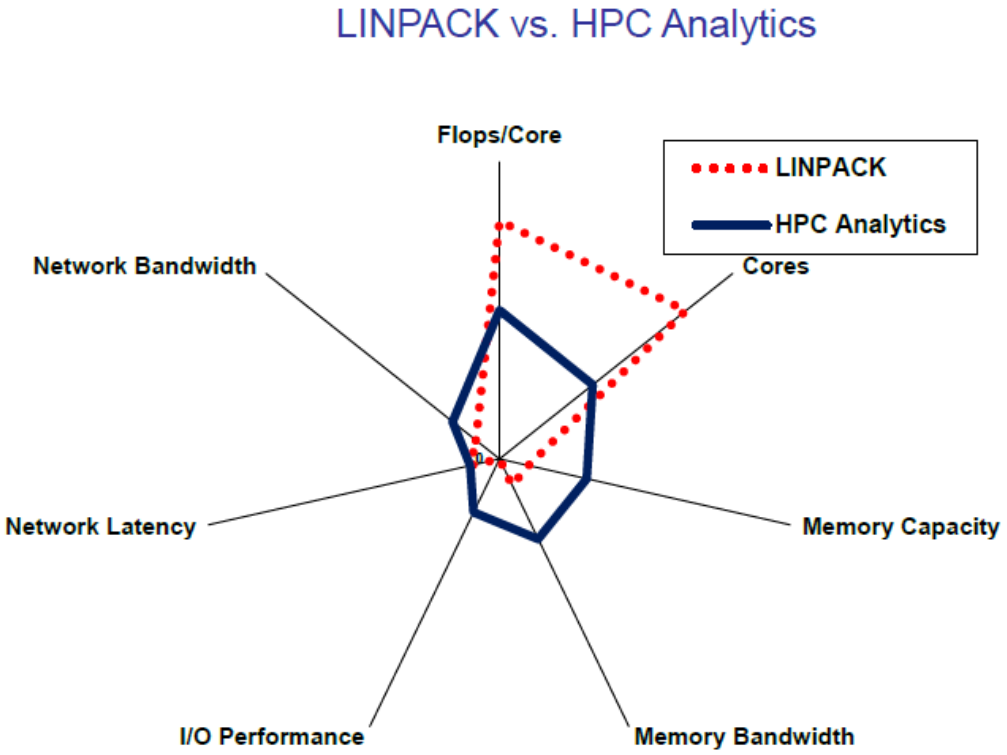

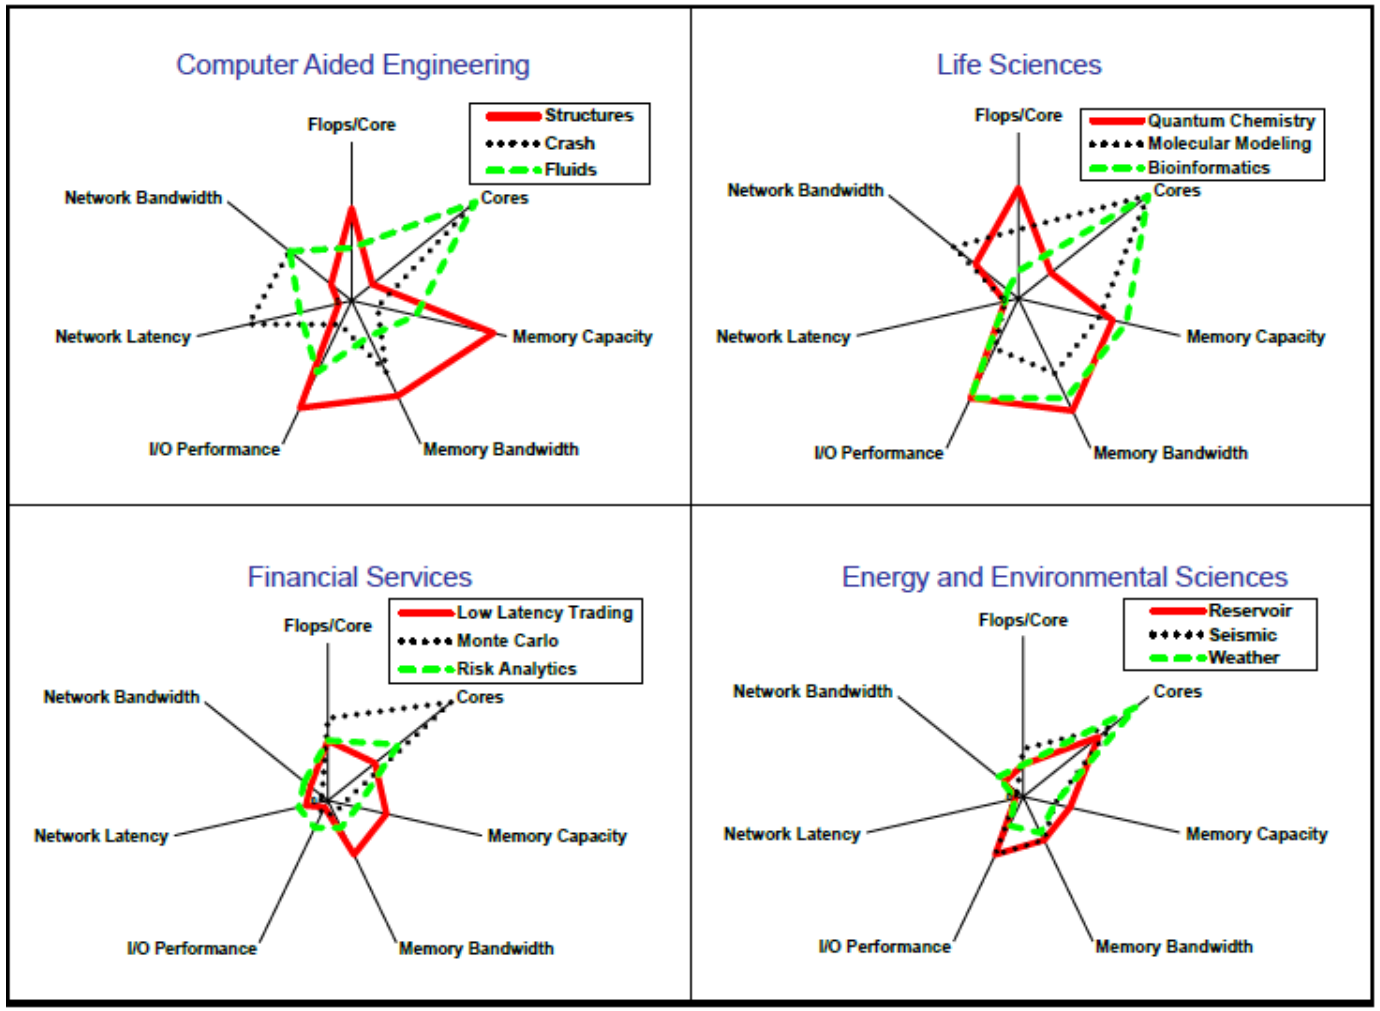

https://www.nextplatform.com/2015/03/19/hpc-scales-out-of-balance-with-cpu-heavy-thinking/

Benchmarks in HPC (2)

Benchmarks covering full workloads

- CORAL 2: https://asc.llnl.gov/coral-2-benchmarks/

Tier-1 Benchmark Information

Scalable Science Benchmarks	Lines of Code	Parallelism			Language				Code Description/Notes
		MPI	OpenMP/ Pthreads	GPU	Fortran	Python	C	C++	
HACC • source • summary • baseline GPU version: **No**	35,000	X	X	X				X	The Hardware Accelerated Cosmology Code (HACC) framework uses N-body techniques to simulate the formation of structure in collisionless fluids under the influence of gravity in an expanding universe. It depends on external FFT library and is typically compute limited achieving 13.92 Petaflops, 69.2% of machine peak on Sequoia.
Nekbone • source • summary • baseline GPU version: **No**	48,000	X		X	X		X		Nekbone is a mini-app derived from the Nek5000 CFD code which is a high order, incompressible Navier- Stokes CFD solver based on the spectral element method. The conjugate gradiant solve is compute intense, contains small messages and frequent allreduces.
QMCPACK • source • summary • baseline GPU version: **Yes**	200.000	X	X	X			X	X	QMCPACK is a many-body ab initio quantum Monte Carlo code for computing the electronic structure of atoms, molecules, and solids. It is written primarily in C++, and its use of template metaprogramming is known to stress compilers. When run in production, the code is memory bandwidth sensitive, while still needing thread efficiency to realize good performance.
LAMMPS • source • summary • baseline GPU version: **Yes**	500,000	X	X	X			X	X	LAMMPS is a classical molecular dynamics code. Performance limiters will depend on the problem chosen and could include, compute, memory bandwidth, network bandwidth, and network latency.

Specialized architectures in HPC

- Hardware architectures in Supercomputing diversify
  - GPUs
  - manycore CPUs
  - FPGAs
- Multiple architectures are combined to heterogenous systems to increase Performance in HPC
- Paderborn is at forefront of FPGA accelerators in HPC

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
3	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 , NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482

Top 4 of the Top500 List (06/2019) (https://www.top500.org/lists/2019/06/)

→ Meaningful Benchmarks for FPGAs are missing

Example: Noctua

- Inaugurated in September 2018
- Technical Description
  - Cray CS500 system
  - 10880 cores in total

  - 256 compute nodes
    - 2x Intel Xeon Gold “Skylake” 6148 (40 cores)
    - 192 GiB memory
  - 16 FPGA nodes
    - 2x Intel Xeon Gold “Skylake” 6148 (40 cores)
    - 192 GiB memory
    - 2x Intel Stratix 10 FPGA

How to choose the best FPGA card for a next system?

FPGA Utilization

- Configurable hardware
  - Customize operations, connections, data reuse

OpenCL

- OpenCL – increase productivity
  - C-like programming language
  - Toolchains available for different FPGA vendors (Intel, Xilinx)
  - Simplifies porting code to different FPGAs

- OpenCL code is used in applications to accelerate compute intensive tasks
  - Which FPGA offers the best performance for these tasks?

The FBench Project Group

- Goals
  - Get familiar with FPGA programming using OpenCL
  - Develop a benchmark suite for FPGA
  - Multiple small benchmarks targeting different performance metrics

- Fields of interest / skills
  - Accelerator architectures and OpenCL
  - Performance Characterization and Comparison
  - Working on remote FPGA machines (Linux, ssh, scripts)

- Supervisors
  - **Christian Plessl**, christian.plessl (at) uni-paderborn.de
  - **Marius Meyer**, marius.meyer (at) uni-paderborn.de
  - **Tobias Kenter**, kenter (at) uni-paderborn.de

Project Goals

- Simple, synthetic benchmarks to cover relevant performance metrics
  - Global memory bandwidth with different access patterns
  - Local memory bandwidth with different data partitions
  - Operations (floating point, integer, logic) using DSPs and LUTs
  - Internal communication on chip, routing limits
  - External communication to other FPGAs
- Benchmark suite that is easy to use
  - Straight forward OpenCL code + documentation
  - Clone repository, type make, get score + explanation
- Optimized benchmark versions for some FPGAs
  - Special variants to perform better than the simple code
- Provide parameters for scaling
  - Peak operations / local memory requires usage of all (~90%) resources

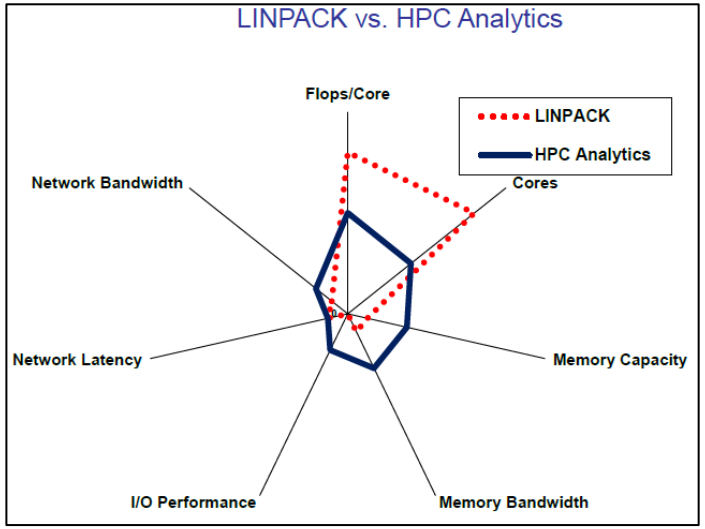

https://www.nextplatform.com/2015/03/19/hpc-scales-out-of-balance-with-cpu-heavy-thinking/

Organization

- Customer role of supervisors for benchmark requirements
  - Can contribute own ideas, but not a completely open project
- Qualification process will require personal initiative
  - Self training with documentation while working on first simple benchmark kernels
  - https://www.intel.com/content/dam/altera-www/global/en_US/pdfs/literature/hb/opencl-sdk/aocl_programming_guide.pdf (208 pages)
  - https://www.intel.com/content/dam/altera-www/global/en_US/pdfs/literature/hb/opencl-sdk/aocl-best-practices-guide.pdf (191 pages)
  - No seminar in first semester
- Mandatory meeting on Friday 9-14 (before HPC lecture)
  - Individual work and meetings on agreement
- Send email with statement of interest + background (see homepage)
  - Additional testing depending on interest, will be announced by August 10

https://cs.uni-paderborn.de/hit/teaching/courses/ws-201920/fbench-defining-and-optimizing-opencl-benchmarks-for-fpgas/