

Democratizing AI Accelerators for HPC Applications: Challenges, Success, and Support

ISC 2024 | May 14, 2024 | Josef Weidendorfer

Collaboration between

BDAI Team (Big Data Artificial Intelligence)

Nicolay Hammer, Juan Durillo, Jophin John, Michael Hoffmann

FC Team (Future Computing)

Josef Weidendorfer, Amir Raoofy, Arjun Parab



- Part of German GCS (Gauss Centre for Supercomputing)
 - Compute proposals to be granted
 - German / European
- Academics in Munich area

~ 2.000
Researchers



100s of projects

- Astrophysics, Particle Physics
- Chemistry / Material Science
- Comp. Fluid Dynamics / Eng.
- Environmental / Life Sciences



LRZ user base is diverse

- **codes often developed by users**
- **porting/tuning is not contributing to science, so often not in focus**

SuperMUC-NG

Top500 - Nov 2018: #8
(Nov 2021: #23)

Lenovo Intel

311,040 cores


Intel Xeon Skylake

26.9 PetaFlops Peak

19.5 PetaFlops Linpack*

719 TeraByte Main Memory

70 PetaByte Disk



Accelerated node architecture

- 2x Intel® Xeon® Platinum 8480L, 56 cores
- 4x Intel® Data Center GPU Max 1550
- 512 GB DDR5 main memory
- Lenovo's SD650-I v3 platform



Distributed asynchronous object storage (DAOS)

- 1 PB capacity
- > 750 GB/s write bandwidth



High speed interconnect

- Mellanox HDR Infiniband
- fat tree topology
- two uplinks per node
- separated from Phase 1



Integration

- Phase 1 accounts and HOME directories
- Phase 1 WORK and SCRATCH filesystems
- DSS volumes available
- direct warm water cooling



Efforts on Future Technologies: BEAST Experimental Environment



BEAST (Bavarian Energy Architecture Software Testbed)


- Evaluate recent hardware technology options for HPC and AI
 - CPUs : x86 (Intel / AMD), ARM (Marvell, Fujitsu, Nvidia GH)
 - GPUs : Nvidia (A-100/H-100), AMD (MI-100/200), Intel (PVC)
 - Special Purpose Accelerators : NextSilicon DataFlow, Cerebras WSE-2
- Other research infrastructure (FPGAs, Smart Network, Quantum Computing)

BEAST (Bavarian Energy Architecture Software Testbed)

- Benchmarking of hardware located at LRZ allows to understand
 - Performance for different domains (vs. vendor claims)
 - Power/Energy efficiency measurements (using LRZ monitoring solution DCDB)
 - Stability of Software Stack
 - Experience from selected friendly users
 - Effectiveness of vendor-provided support on hardware failures

Evaluation of Cerebras WSE-2 as part of BEAST

In Contact with 13 “Friendly” User Groups



“Unconventional Usage“ of Accelerators

Can we support HPC via AI Accelerators?

Why does anybody want to use an exotic AI HW for HPC?

- Exceptional On-Chip Memory Capability
 - Size
 - Bandwidth
- Huge number of cores on a single chip promise low-latency synchronization

- Unique architecture of WSE-2
 - 850000 cores
 - 40GBs on-chip SRAM
 - 20 PB/s memory bandwidth
 - Very high NoC bandwidth

Case Study (In Cooperation with TUM)



Master Thesis

- “Implementation and Evaluation of Matrix Profile Algorithms on the Cerebras Wafer-Scale Engine”
- Time Series Mining (Similarity Indexing)

- Matrix profile computations are usually performed on CPUs and GPUs using SCAMP library which is known to be fast
- Can we port SCAMP to CS-2 WSE?

- Experiments were conducted on the CS-2 of the EIDF (EPCC)

Case Study - Preliminary Take Away Lessons

- Cerebras's SDK and Documentation have matured over the study's time allowing to successfully port the SCAMP Kernel to CS-2
 - Functional port successful
 - No optimizations yet done about
 - on-chip data pipelining
 - data transfers to device
 - Challenge to make best use of on-chip per-core memory (both needed by code and data)

Is Cerebras WSE Usable for HPC?



Usecase Study

- First efforts promising
- More research required
- Default software stack may be difficult for HPC usage

General Perspective

- Ongoing evolution (CS-2 to CS-3)
- Expanded library and customer support enhances usability
- Continuous feature additions to the SDK