SC24 Atlanta, GA | hpc creates.

Trip Report

16.12.2024

Jolly Chen – CERN EP-SFT / University of Twente

Overview

- November 17-22 · Atlanta, GA
- Over **18.000** attendees and **494** exhibitors
- **90** research papers from **18** countries, **44** workshops, **36** tutorials, **89** birds of feathers, **77** research posters, and **16** doctoral showcase posters
    - 5 Contributions on Julia: 1 BoF, 2 workshop papers, 1 poster, 1 tutorial

Overview

- **Themes**
  - Algorithms
  - Applications
  - Architectures and networks
  - Clouds and distributed computing
  - Data analytics, visualization, and storage
  - Machine learning and HPC
  - Performance
  - Programming systems
  - System software
  - State of the practice in large-scale deployment and integration
  - Education and inclusivity

My Contributions

- Presenter in tutorial on “Hands-On HPC and AI Application Development Using C++ and SYCL”
- Panel member in Birds of a Feather session on “Khronos SYCL: Heterogeneous Programming with Open Standards”
  - Presented the use cases of SYCL in ATLAS, CMS, and ROOT

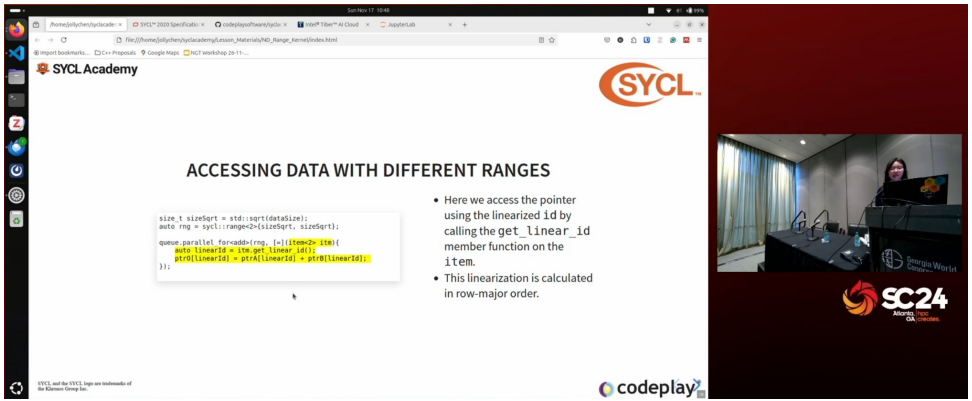

CERN Easter Eggs

Paper

- Reshaping High-Energy Physics Applications for Near-Interactive Execution Using TaskVine
  - **Author** Barry Sly-Delgado, Ben Tovar, Jin Zhou, and Douglas Thain – University of Notre Dame

Poster

- Predicting Dataset Popularity for Improved Distributed Content Caching in High Energy Physics
  - **Author** Malavikha Sudarshan – University of California, Berkeley

CERN Easter Eggs

Workshops

- Optimising Science Workflows with On-Demand Machine Learning Inference on Perlmutter Supercomputer
  - **Workshop** Fourth Combined Workshop on Interactive and Urgent HPC
  - **Author** Andrew Naylor – NERSC
- AI Surrogate Model for Distributed Computing Workloads
  - **Workshop** AI4S: 5th Workshop on Artificial Intelligence and Machine Learning for Scientific Applications
  - **Authors** David K. Park, Yihui Ren, Ozgur O. Kilic, Sairam Sri Vatsavai, Tasnuva Chowdhury, Tadashi Maeno, Paul Nilsson, Shinjae Yoo, Alexei Klimentov, Adolfy Hoisie (Brookhaven National Laboratory), Shengyu Feng, Jaehyung Kim, Yiming Yang (Carnegie Mellon University), Korchuganova, Joseph Boudreau, Raees A. Khan (University of Pittsburgh), Klasky, Norbert Podhorszki, Frédéric Suter (Oak Ridge National Laboratory (ORNL), Ingrid Martinez Outschoorn (University of Massachusetts, Amherst), Wei Yang (SLAC National Accelerator Laboratory)

CERN Easter Eggs

ACM Student Research Competition

- Comparing Cache Utilization Trends for Regional Scientific Caches with Transfer Learning Models
  - **Author** Erica Wang – California Institute of Technology, Lawrence Berkeley National Laboratory (LBNL)

Found at the KEK exhibition booth

Mentioned at the NVIDIA booth

(Personal) Highlights

Invited talk Mapping Irregular Computations to Accelerator-Based Exascale Systems

- **Speaker** Katherine Yelick – University of California, Lawrence Berkeley National Laboratory (LBNL)
- 10 ways to **waste** an exascale system
  - Embrace communication
  - Do not overlap communication and computation
  - Ignore arithmetic innovations
  - Ignore historic parallel algorithm work
  - Ignore/do not look too hard for spatial locality
  - Ignore/do not look too hard for temporal locality
  - Use serial container data structures (STL)
  - Ignore load imbalance
  - (Rigidly) choose between task & data parallelism
  - Let the GPU rule

(Personal) Highlights

Workshop Paper Testing GPU Numerics: Finding Numerical Differences Between NVIDIA and AMD GPUs

- **Workshop** Second International Workshop on HPC Testing and Evaluation of Systems, Tools, and Software (HPCTESTS 2024)
- **Authors** Anwar Hossain Zahid (Iowa State University), Ignacio Laguna (Lawrence Livermore National Laboratory (LLNL)), and Wei Le (Iowa State University)
- Automatically generated short numerical tests with Varity
- **652,600** experimental runs on **NVIDIA V100** and **AMD MI-250X**
- Used Hipify to convert CUDA programs into HIP

(Personal) Highlights

TABLE IV
SUMMARY OF EXPERIMENTAL RESULTS

Metric	FP64	FP64 with HIPIFY	FP32
Total Programs	3,540	3,540	2,840
Total Runs per Option per Compiler	24,750	24,750	15,760
Total Runs per Option	49,500	49,500	31,520
Total Runs	**247,500**	**247,500**	**157,600**
Runs on NVCC	123,750	123,750	78,800
Runs on HIPCC	123,750	123,750	78,800
Total Discrepancies	**2,426**	**2,716**	**14,188**
Total Discrepancies (% of Total Runs)	*0.98%*	*1.10%*	*9.00%*

Significantly more discrepancies with FP32 compared to FP64

(Personal) Highlights

Other interesting results:

```
Expression: fmod(1.5917195493481116e+289, 1.5793E
    -307)

Output:
nvcc -O0: 1.4424471839615770678 0e-307
hipcc -O0: 7.1923082856620736002 9e-309
```

```
Expression: ceil(1.5955E-125)

Output:
nvcc -O0:  0
hipcc -O0: 1
```

(Personal) Highlights

#	Site	Manufacturer	TOP10 Computer of the TOP500	Country	Cores	Rmax [Pflops]	Power [MW]
1	Lawrence Livermore National Laboratory	HPE	El Capitan HPE Cray EX255a, AMD EPYC 24C 1.8GHz, Instinct MI300A, Slingshot-11	USA	11,039,616	1,742	29.6
2	Oak Ridge National Laboratory	HPE	Frontier HPE Cray EX235a, AMD EPYC 64C 2.0GHz, Instinct MI250X, Slingshot-11	USA	9,066,176	1,353	24.6
3	Argonne National Laboratory	Intel	Aurora HPE Cray EX/Intel Exascale Compute Blade, Xeon Max 9470, Data Center GPU Max, Slingshot-11	USA	4,742,808	1,012	38.7
4	Microsoft Azure	Microsoft	Eagle Microsoft NDv5, Xeon Platinum 8480C, NVIDIA H100, Infiniband NDR	USA	1,123,200	561.2	
5	Eni S.p.A. Center for Computational Science	HPE	HPC6 HPE Cray EX235a, AMD EPYC 64C 2.0GHz, Instinct MI250X, Slingshot-11	Italy	3,143,520	477.9	8.5
6	RIKEN Center for Computational Science	Fujitsu	Fugaku Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D	Japan	7,630,848	442.0	29.9
7	Swiss National Supercomputing Centre (CSCS)	HPE	Alps HPE Cray EX254n, NVIDIA Grace 72C 3.1GHz, GH200, Slingshot-11	Switzerland	2,121,600	434.9	7.1
8	EuroHPC / CSC	HPE	LUMI HPE Cray EX235a, AMD EPYC 64C 2.0GHz, Instinct MI250X, Slingshot-11	Finland	2,752,704	379.7	7.1
9	EuroHPC / CINECA	EVIDEN	Leonardo Atos BullSequana XH2000, Xeon 32C 2.6GHz, NVIDIA A100, HDR Infiniband	Italy	1,824,768	241.2	7.5
10	Lawrence Livermore National Laboratory	HPE	Tuolumne HPE Cray EX255a, AMD EPYC 24C 1.8GHz, Instinct MI300A, Slingshot-11	USA	1,161,216	208.1	3.4

BOF Top500 Supercomputers - Erich Stromaier (Top500)

(Personal) Highlights

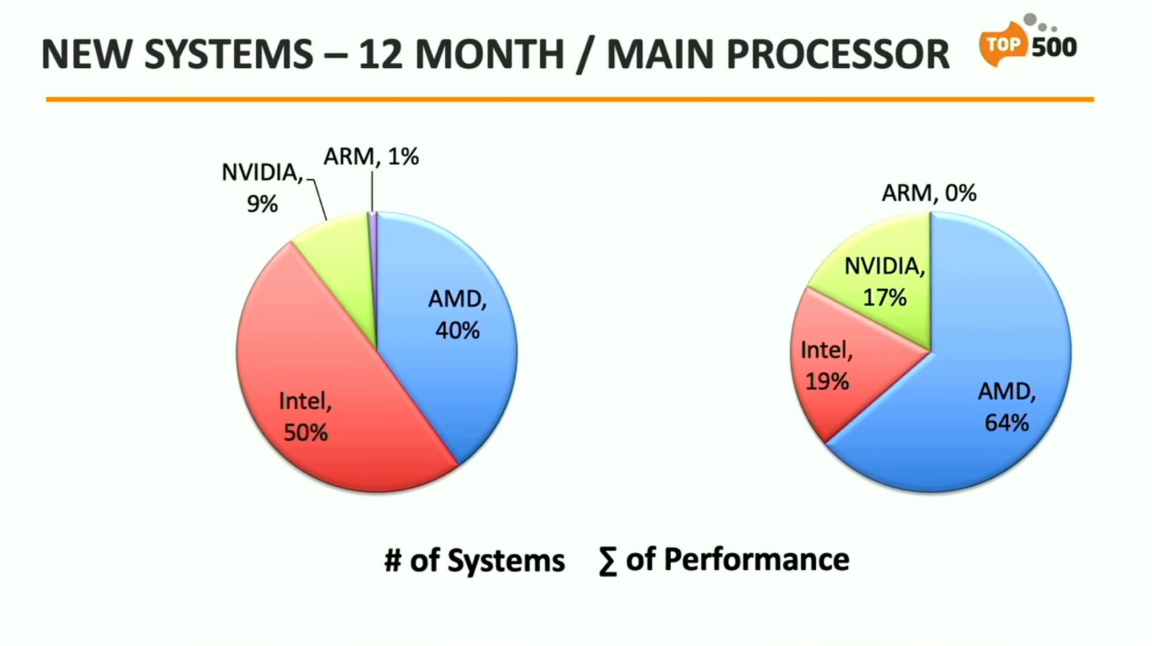

BOF Top500 Supercomputers

(Personal) Highlights

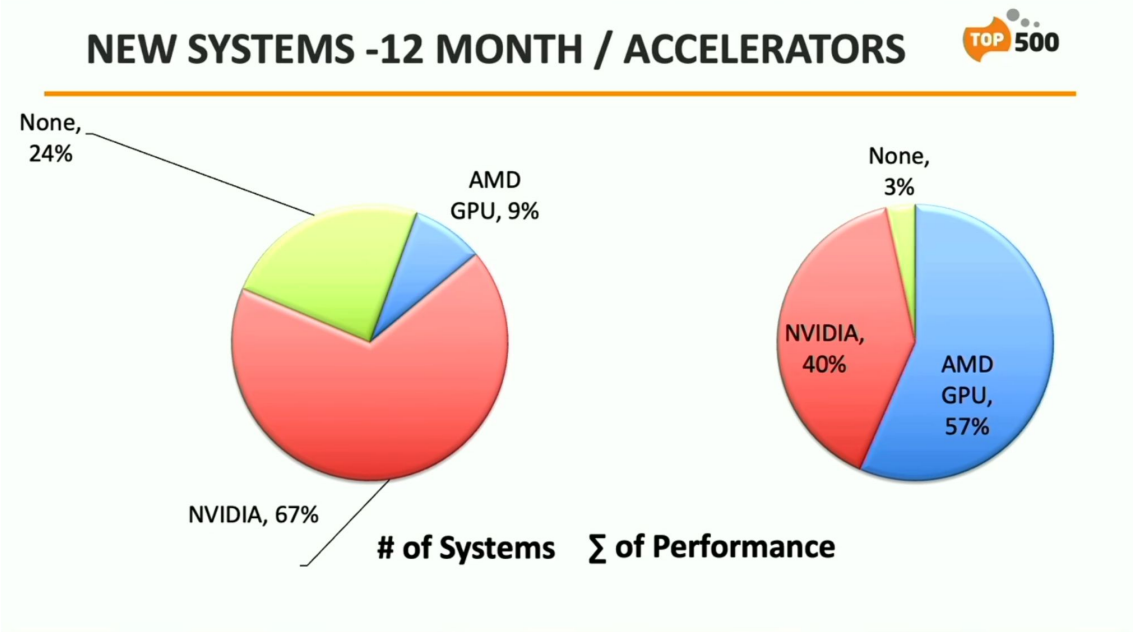

BOF Top500 Supercomputers

(Personal) Highlights

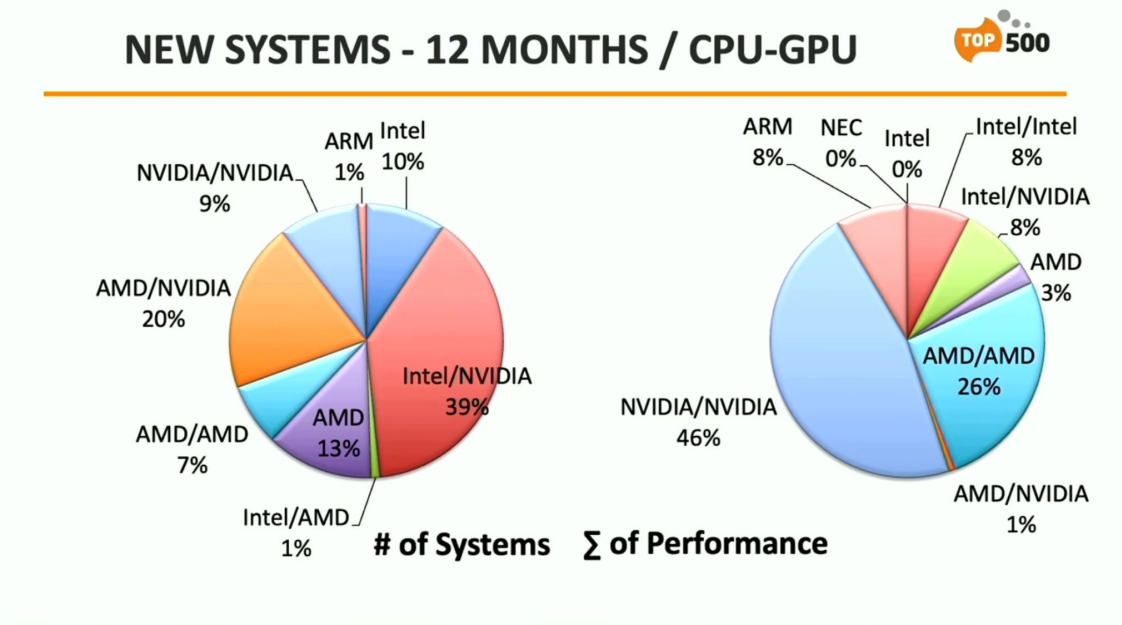

BOF Top500 Supercomputers

Conclusion

- *Many* things happening in parallel
- Some contributions about CERN, but not by CERN
- Things that I felt were (unexpectedly) mentioned a lot: Grace Hopper GPUs, digital twins, LLMs, quantum, Fortran, SYCL, …