UCF Unified Communication Framework

Welcome to 7th UCF Virtual Workshop

Pavel (Pasha) Shamis & Gilad Shainer

UCF Workshop, 2021

This is an open, public standards setting discussion and development meeting of UCF. The discussions that take place during this meeting are intended to be open to the general public and all work product derived from this meeting shall be made widely and freely available to the public. All information including exchange of technical information shall take place during open sessions of this meeting and UCF will not sponsor or support any closed or private working group, standards setting or development sessions that may take place during this meeting. Your participation in any non-public interactions or settings during this meeting are outside the scope of UCF's intended open-public meeting format.

Housekeeping Notes

- Workshop page with up-to-date schedule
  - https://github.com/openucx/ucx/wiki/7th-Annual-UCF-Workshop-and-Annual-Meeting-2021
  - Zoom link - the same link will work all 3 days
    - https://armltd.zoom.us/j/99916624100?pwd=R2VUa2hicEdrM2RnRm1DQjN3clNmUT09
- Session Chairs
  - Pavel (Pasha) Shamis, Oscar Hernandez, Matt Bakker
- All presentations and discussions are video recorded
  - Zoom explicitly asks participants for a consent
  - Videos and slides will be posted online
- Etiquette
  - Please stay on mute unless you are presenting ☺
  - Presenters are encouraged to turn video on
  - If you participate in a discussion or ask a question, it is encouraged to turn the video on
  - If you have questions during a presentation, please wait until end of presentation or use zoom “raise the hand” to signal that you have questions. You also can use Zoom chat to ask questions.
  - You can use Zoom whiteboard and screen sharing for discussion sessions

Agenda – Day 1

Date	Time	Topic	Speaker/Moderator
11/30	08:00-08:15	▶ Opening Remarks and UCF	▶ Gilad Shainer, NVIDIA
	08:15-09:00	▶ Accelerating recommendation model training using ByteCCL and UCX	▶ Haibin Lin*, Bytedance Inc. ▶ Mikhail Brinskii, NVIDIA ▶ Yimin Jiang, Bytedance Inc. ▶ Yulu Jia, Bytedance Inc. ▶ Chengyu Dai, Bytedance Inc. ▶ Yibo Zhu, Bytedance Inc.
	09:00-09:30	▶ UCX on Azure HPC/AI Clusters	▶ Jithin Jose, Microsoft
	09:30-10:00	▶ MPICH + UCX: State of the Union	▶ Ken Raffenetti, Argonne National Laboratory
	10:00-10:30	Break	
	10:30-11:00	▶ UCX.jl -- Feature rich UCX bindings for Julia	▶ Valentin Churavy, MIT
	11:00-11:30	▶ Go bindings for UCX	▶ Peter Rudenko, NVIDIA
	11:30-12:00	▶ UCX-Py: Stability and Feature Updates	▶ Peter Entschev, NVIDIA
	12:00-12:30	Break	
	12:30-13:00	▶ OpenSHMEM and Rust	▶ Tony Curtis*, Stony Brook University ▶ Rebecca Hassett, Stony Brook University
	13:00-13:30	▶ Towards Cost-Effective and Elastic Cloud Database Deployment via Memory Disaggregation	▶ Cheng Li, University of Science and Technology of China
	13:30-13:45	Adjourn	

Agenda – Day 2

12/01	08:00-08:45	Opening Remarks and UCX	▶ Pavel Shamis (Pasha), Arm
	08:45-09:00	▶ Porting UCX for Tofu-D interconnect of Fugaku	▶ Yutaka Watanabe*, University of Tsukuba ▶ Mitsuhisa Sato, RIKEN ▶ Miwako Tsuji, RIKEN ▶ Hitoshi Murai, RIKEN ▶ Taisuke Boku, University of Tsukuba
	09:00-10:00	▶ UCP Active Messages	▶ Mikhail Brinskii*, NVIDIA ▶ Yossi Itigin, NVIDIA
	10:00-10:30	Break	
	10:30-12:00	▶ UCX GPU support	▶ Yossi Itigin*, NVIDIA ▶ Akshay Venkatesh, NVIDIA ▶ Devendar Bureddy, NVIDIA
	12:00-12:30	Break	
	12:30-13:00	▶ UCX on AWS: Adding support for Amazon EFA in UCX	▶ Hessam Mirsadeghi*, NVIDIA ▶ Akshay Venkatesh, NVIDIA ▶ Jim Dinan, NVIDIA ▶ Sreeram Potluri, NVIDIA
	13:00-13:30	▶ rdma-core update	▶ Jason Gunthorpe, NVIDIA
	13:30-14:00	▶ Congestion Control for Large Scale RoCEv2 Deployments	▶ Hemal Shah*, Broadcom ▶ Moshe Voloshin, Broadcom
	14:00-14:15	Adjourn	

Agenda – Day 3

12/02	08:00-08:15	▶ Opening Remarks	▶ Pavel Shamis (Pasha), Arm
	8:15-09:30	▶ Unified Collective Communication (UCC) State of the Union 2021	▶ Manjunath Gorentla Venkata, NVIDIA
	09:30-10:00	▶ High Performance Compute Availability (HPCA) Benchmark Project for Smart Networks	▶ Geoffroy Vallee*, NVIDIA ▶ Richard Graham, NVIDIA ▶ Steve Poole, Los Alamos National Laboratory
	10:00-10:30	Break	
	10:30-11:00	▶ Cloud-Native Supercomputing Performance isolation	▶ Gilad Shainer*, NVIDIA ▶ Richard Graham, NVIDIA ▶ Jithin Jose, Microsoft
	11:00-12:00	▶ ifunc: UCX Programming Interface for Remote Function Injection and Invocation	▶ Wenbin Lu*, Stony Brook University ▶ Luis E. Peña, Arm Research
	12:00-12:30	Break	
	12:30-13:00	▶ Remote OpenMP Offloading with UCX	▶ Atmn Patel*, University of Waterloo ▶ Johannes Doerfert, Argonne National Laboratory
	13:00-13:30	▶ From naive to smart: leveraging offloaded capabilities to enable intelligent NICs	▶ Whit Schonbein, Sandia National Laboratories
	13:30-14:00	▶ Using Data Processing Units to manage large sets of small files	▶ Matthew Baker, Oak Ridge National Laboratory
	14:00-14:15	Adjourn	

Acknowledgments

- UCF board and organizing committee
- Special thanks to Oscar Hernandez, Matt Bakker, Qingchun Song for all the help with workshop organization

UCF Unified Communication Framework

UCF – State of the Union

Gilad Shainer

Unified Communication Framework (UCF) Consortium

MISSION: Collaboration between industry, laboratories, and academia to create production grade communication frameworks and open standards for data centric, ML/AI, and high-performance applications

- Projects & Working Groups
  - **UCX – Unified Communication X – www.openucx.org**
  - **UCC – Collective Library**
  - **OpenSNAPI – Smart networks Project**
  - **OpenHPCA – Benchmarking Effort**

- Board members
  - **Jeff Kuehn**, UCF Chairman (Los Alamos National Laboratory)
  - **Gilad Shainer**, UCF President (Nvidia)
  - **Pavel Shamis**, UCF Treasurer (Arm)
  - **Brad Benton**, Board Member (AMD)
  - **Yanfei Guo**, Board Member (Argonne National Laboratory)
  - **Perry Schmidt**, Board Member (IBM)
  - **Dhabaleswar K. (DK) Panda**, Board Member (Ohio State University)
  - **Steve Poole**, Board Member (Open Source Software Solutions)

Join https://www.ucfconsortium.org or info@ucfconsortium.org

Why UCX ?

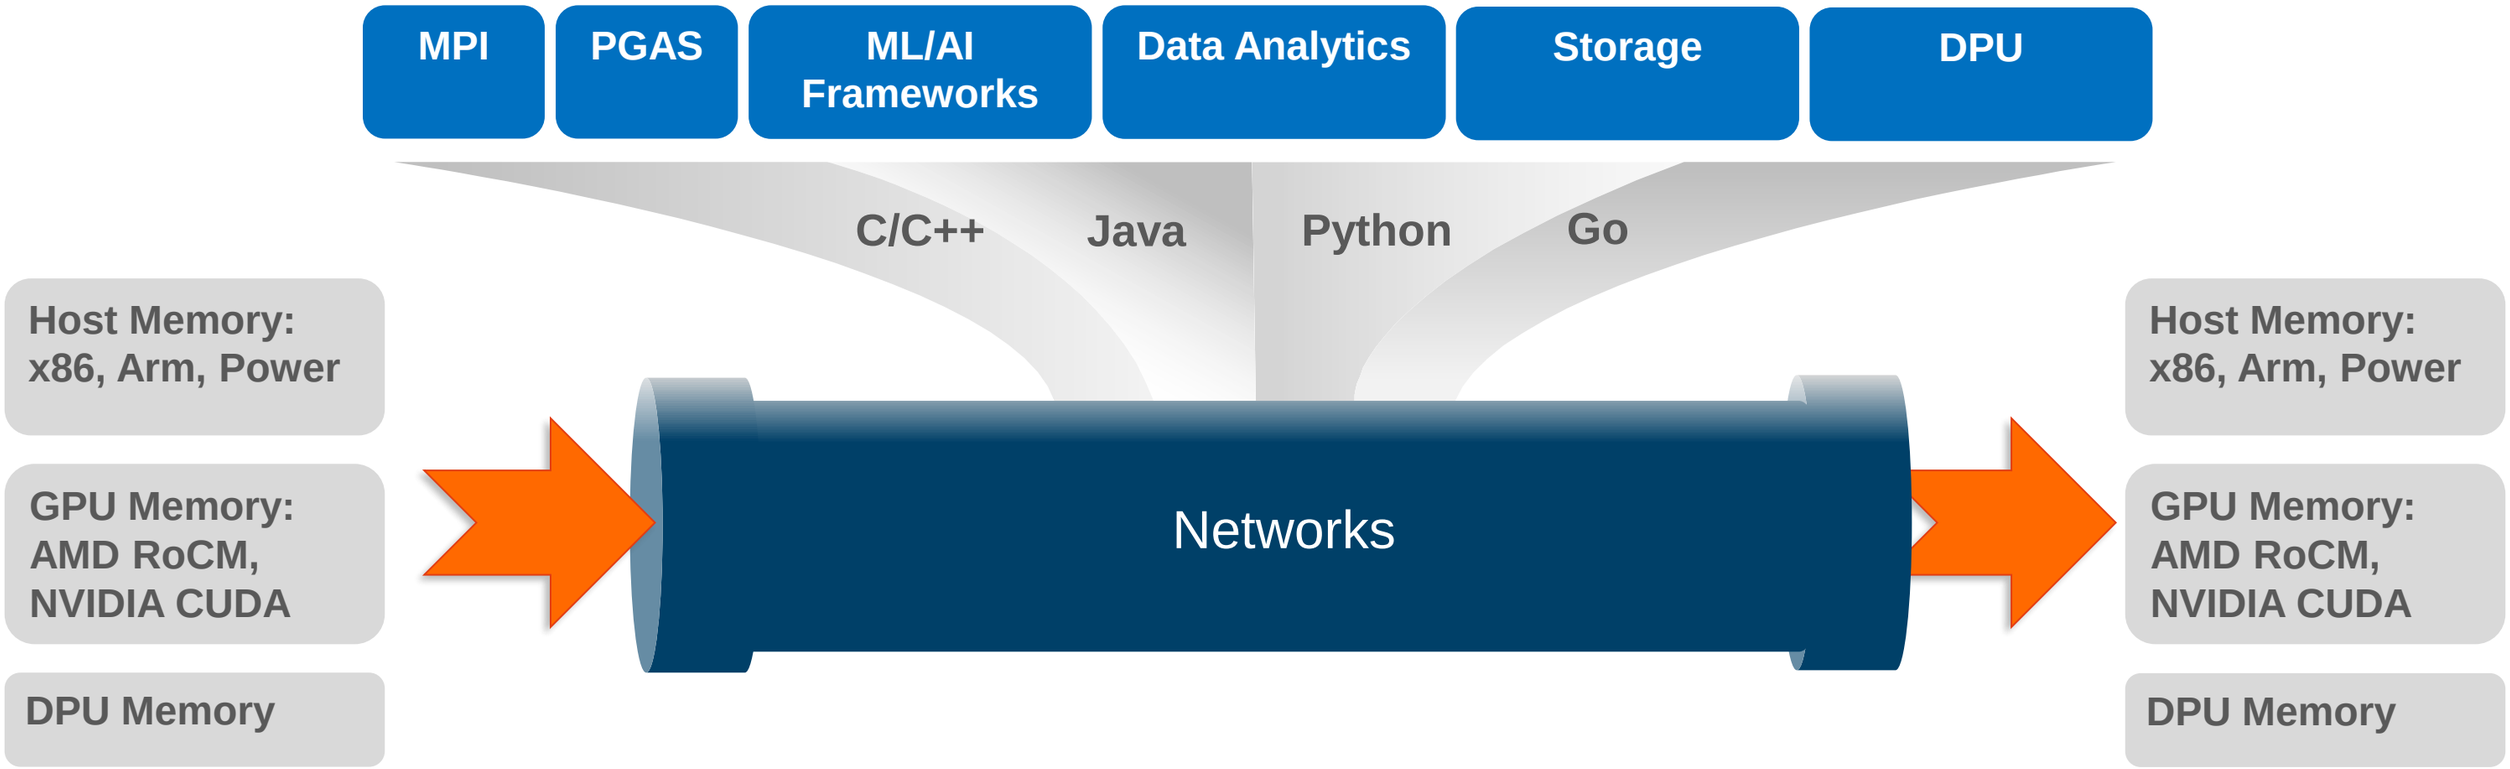

High-Performance Universal Data Mover

UCX Users

- MPI implementations: MPICH, Open MPI, NVIDIA HPC-X, Huawei MPI
- PGAS: GasNET
- OpenSHMEM: OSSS SHMEM, Sandia SHMEM, Open MPI SHMEM
- Charm++
- RAPIDS / DASK
- Spark UCX
- NVIDIA NCCL

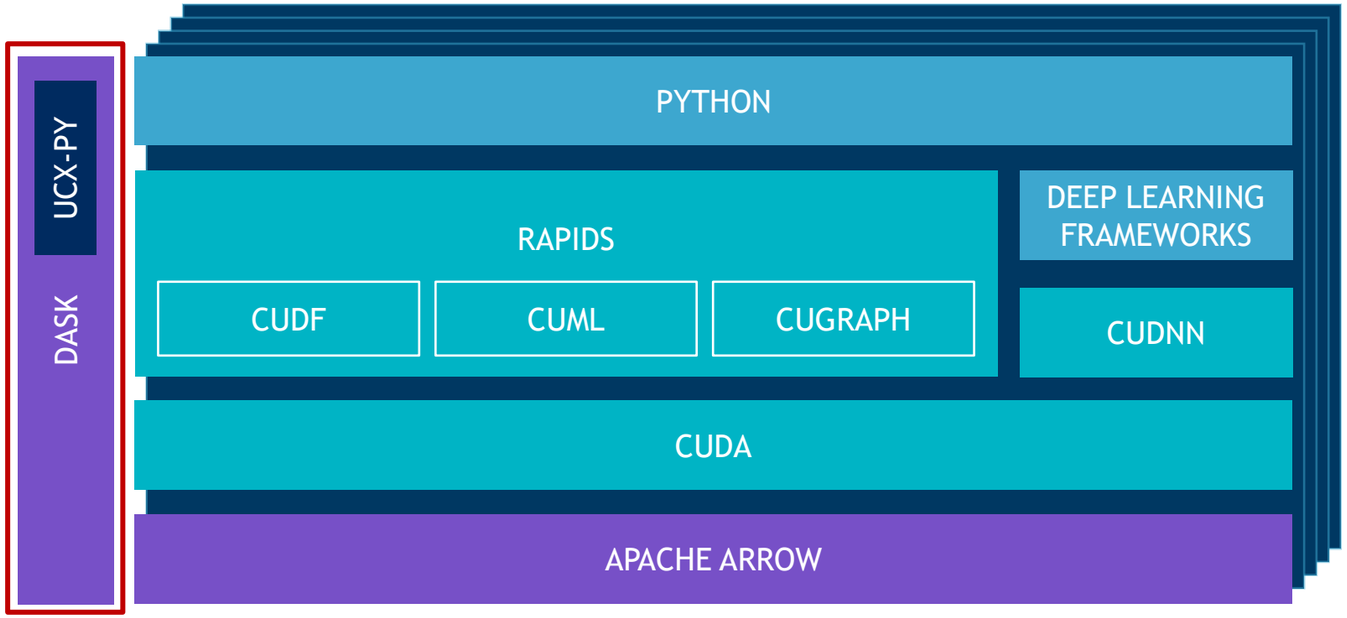

Diagram courtesy of NVIDIA

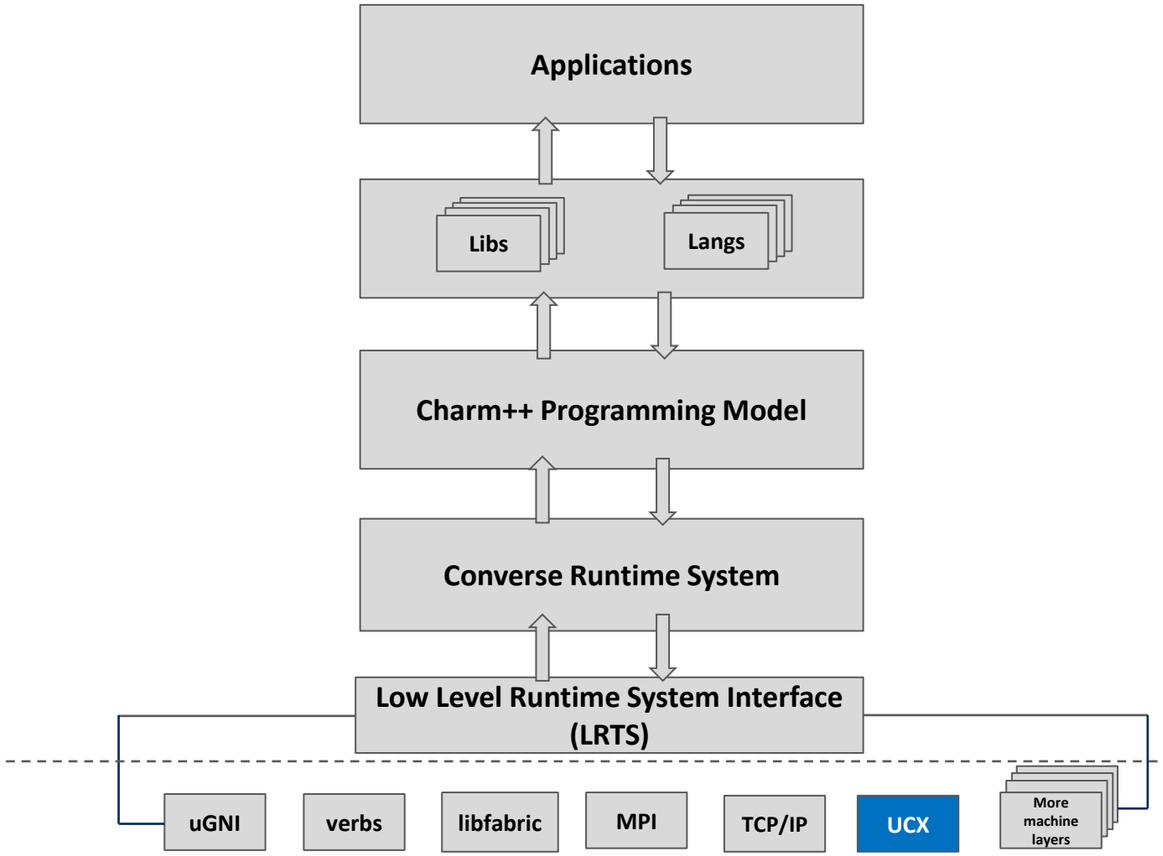

Diagram courtesy of Nitin Bhat @ Charmworks Inc

Content

UCX – Unified Communication X
Web https://www.openucx.org
Git https://github.com/openucx/ucx
Docs https://openucx.readthedocs.io
Mailing list https://elist.ornl.gov/mailman/listinfo/ucx-group

UCC - Collective Communication API
Web https://www.ucfconsortium.org/projects/ucc/
Git https://github.com/openucx/ucc
Git https://github.com/openucx/ucc_spe
Git https://github.com/openucx/torch-ucc

High Performance Compute Availability (HPCA) Benchmark
Web https://ucfconsortium.org/projects/hpca-benchmark/
GIT https://github.com/openucx/openhpca

API for Smart Network (DPU) programmability
Web https://www.ucfconsortium.org/projects/opensnapi/
GIT https://github.com/openucx/ucx-two-chains

MPICH community update

UCF Panel Q&A

Open-source framework for high-performance networks

UCF Unified Communication Framework

Thank You