UCF Unified Communication Framework

Welcome to UCF Virtual Workshop

Pavel Shamis & Gilad Shainer

UCF Workshop, 2020

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/UCF-Virtual-Workshop-2020
- Moderators
  - Pavel Shamis, Cydney Stevens, Brian Sparks
- All presentations and discussions are video recorded
  - Zoom explicitly asks participants for a consent
- 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	08:00-09:00	▶ UCF State of the Union	▶ Gilad Shainer, Nvidia ▶ Pavel Shamis (Pasha), Arm
	09:00-10:00	▶ GPU memory support	▶ Yossi Itigin, Nvidia
	10:00-10:30	▶ MPICH/UCX Update	▶ Ken Raffenetti ,Argonne National Laboratory
	10:30-11:15	▶ RDMA-CORE: DMA-BUF based GPU RDMA Support	▶ Jianxin Xiong, Intel
	11:15-12:15	▶ UCX for Apache Spark	▶ Peter Rudenko, Nvidia
	12:15-13:15	▶ UCX Python - Dask/RAPIDS	▶ Ben Zaitlen, NVIDIA ▶ Peter Entschev, NVIDIA ▶ Matthew Baker, ORNL

Agenda – Day 2

12/01	08:00-08:40	UCF - Future directions	▶ Steve Poole, Los Alamos National Laboratory
	08:40-09:40	▶ UCP Protocols v2	▶ Yossi Itigin, Nvidia
	09:40-10:40	▶ UCP Active messages API	▶ Mikhail Brinskii, Nvidia
	10:40-11:40	▶ UCX development in Huawei	▶ Alex Margolin, HPC software architect and team leader, Huawei
	11:40-12:20	▶ Open Smart NIC API - State of the Union	▶ Steve Poole, Los Alamos National Laboratory

Agenda – Day 3

12/02	08:00-09:00	▶ BlazingSQL with UCX	▶ Rodrigo Aramburu ▶ Felipe Aramburu, BlazingSQL
	09:00-10:00	▶ Charm++ with UCX	▶ Nitin Bhat, Charmworks ▶ Jaemin Choi, University of Illinois Urbana-Champaign
	10:00-10:30	▶ ROCM support in UCX: Status and Roadmap	▶ Sourav Chakraborty, AMD
	10:40-11:40	▶ UCX counters in Score-P and Vampir	▶ Shuki Zanyovka, HPC and Networking architect at Huawei
	11:40-12:40	Unified Communication Datatypes - State of the Union	Pavan Balaji, Argonne National Laboratory
	12:40-13:00	▶ Arm IP building blocks and standards for SmartNIC	▶ Kshitij Sudan, Arm

Agenda – Day 4

12/03	08:00-09:00	▶ UCC: Design and Implementation of Next Generation Collectives Library	▶ Manjunath Gorentla Venkata, Nvidia
	09:00-09:30	▶ One-to-many UCT transports, part I: Shared-memory	▶ Alex Margolin, HPC software architect and team leader, Huawei
	09:30-10:00	▶ One-to-many UCT transports, part II: Multicast	▶ Morad Horany, HPC software developer, Huawei
	10:00-11:00	▶ Until UCC is available - UCG status update	▶ Alex Margolin, HPC software architect and team leader, Huawei
	11:00-11:45	▶ RDMA-CORE Linux kernel and user space updates	▶ Jason Gunthorpe, Nvidia
	11:45-12:45	▶ Scaling Facebook's Deep Learning Recommender Model (DLRM) with UCC/XCCL	▶ Josh Ladd, Nvidia Srinivas, Facebook
	12:45-13:30	▶ Open Smart NIC API - OpenSHMEM I/O Extensions for Fine-grained Access to Persistent Memory Storage	▶ Megan Grodowitz, Arm

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**
  - SparkUCX – www.sparkucx.org
  - OpenSNAPI – Smart NIC Project
  - UCC – Collective Library
  - UCD – Advanced Datatype Engine
  - HPCA Benchmark – 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)
  - **Pavan Balaji**, Board Member (Argonne National Laboratory)
  - **Sameh Sharkawi**, 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

Unified Communication X (UCX)

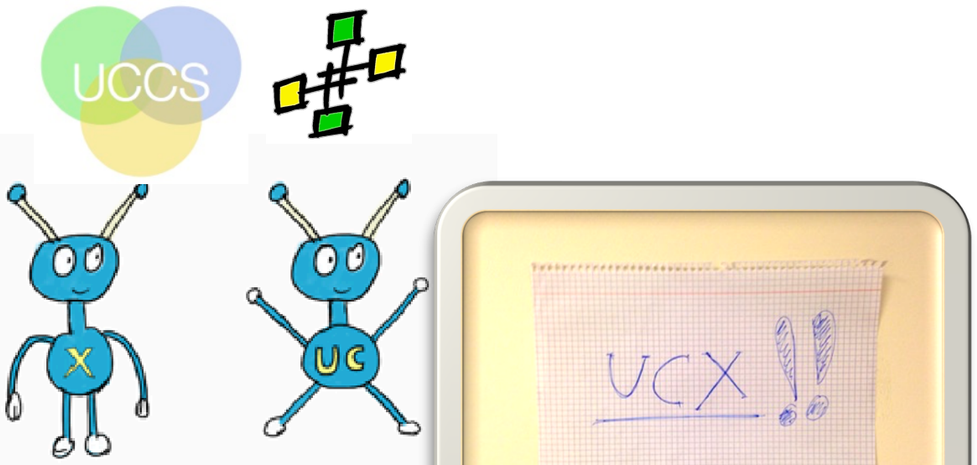

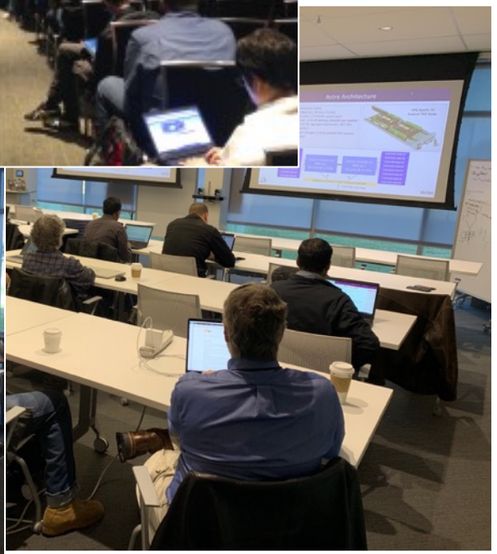

https://www.hpcwire.com/2018/09/17/ucf-ucx-and-a-car-ride-on-the-road-to-exascale/

UCX – History

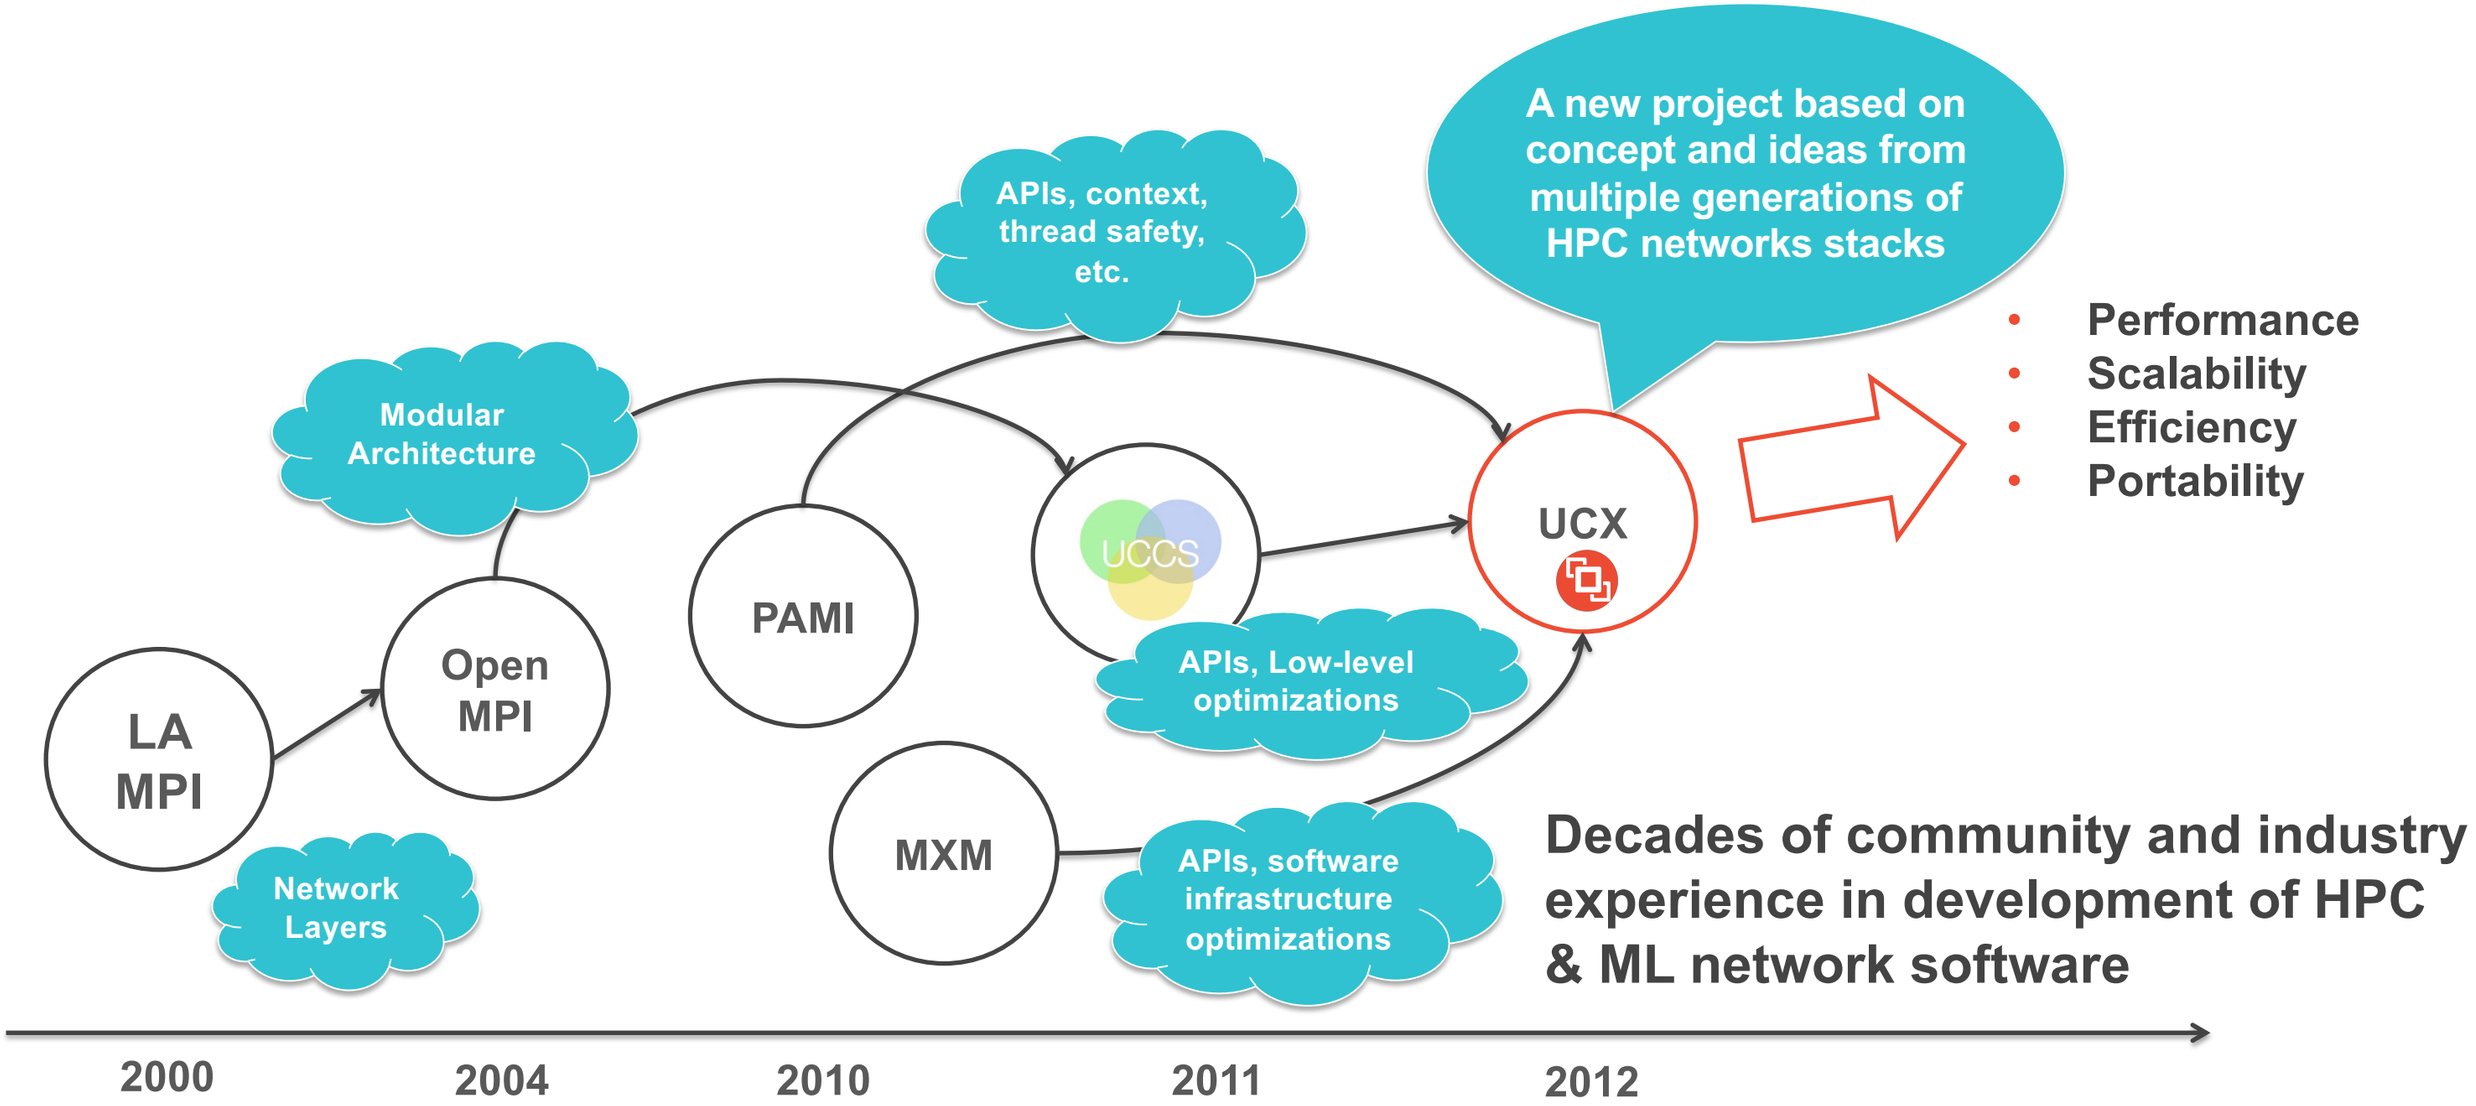

Why UCX ?

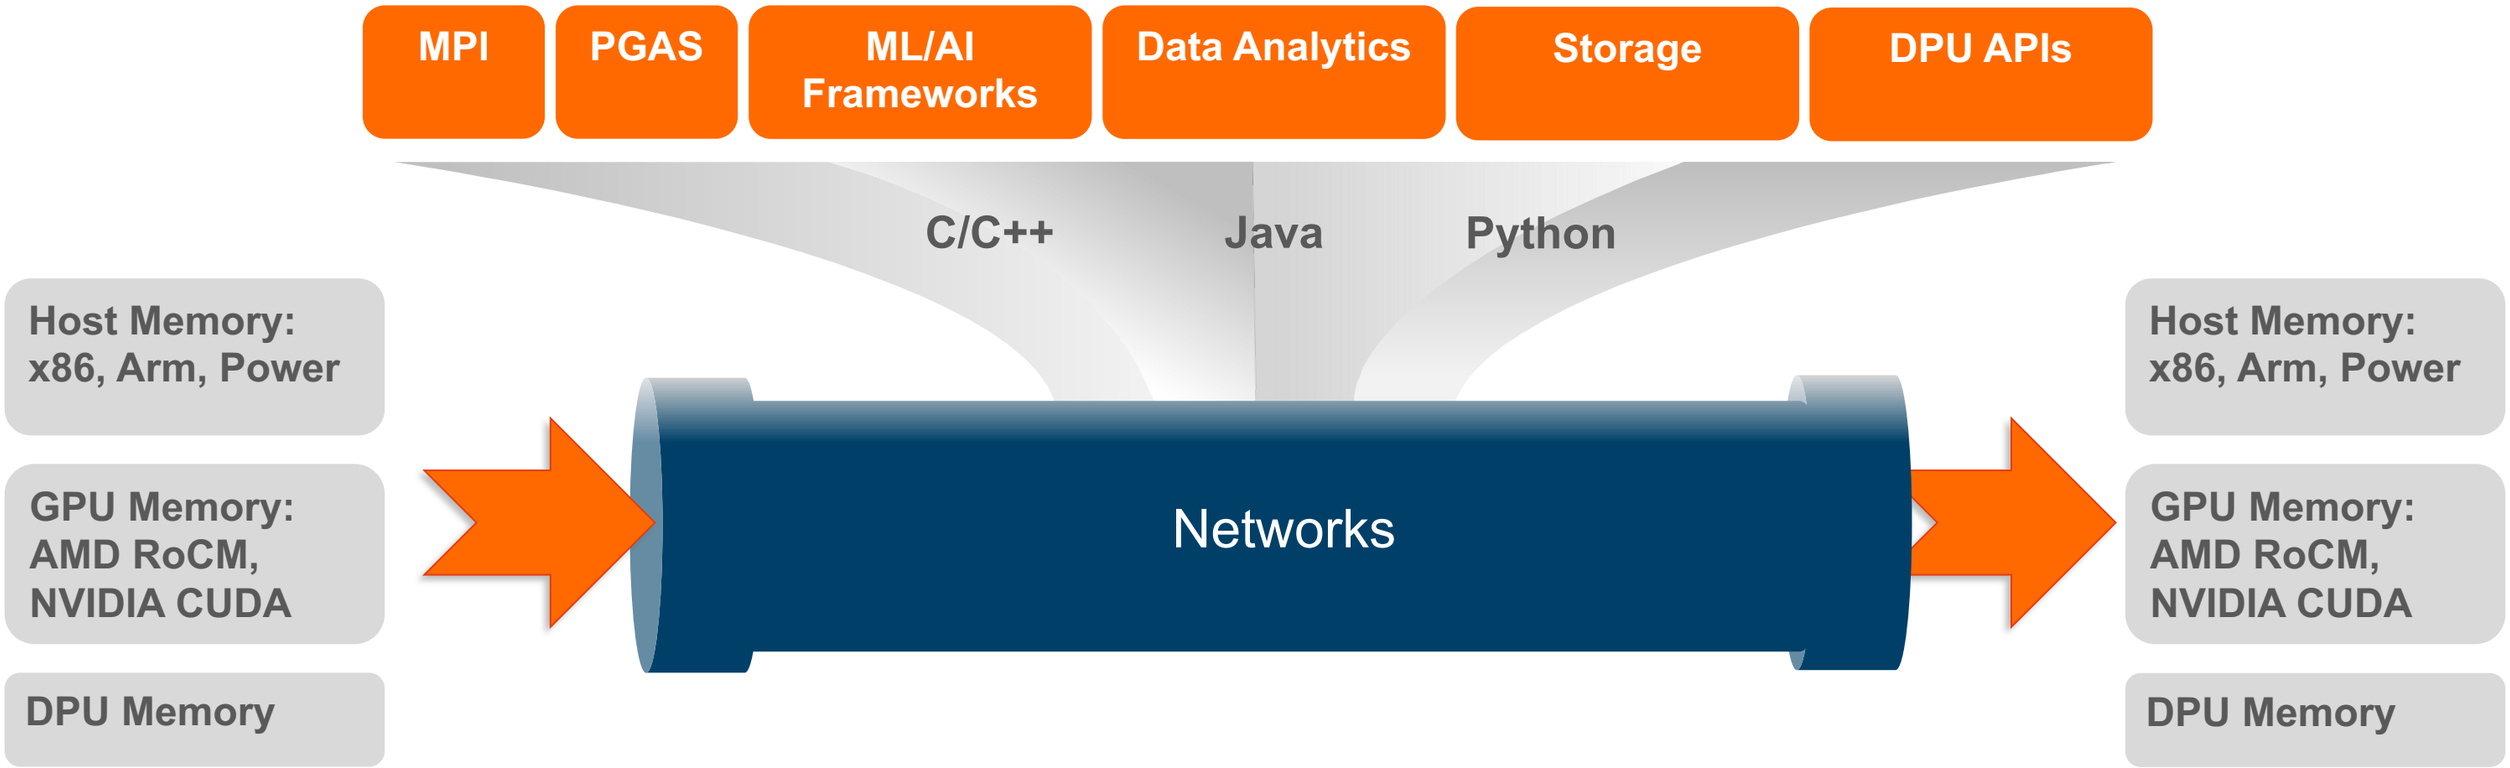

High-Performance Universal Data Mover

2019 R&D 100

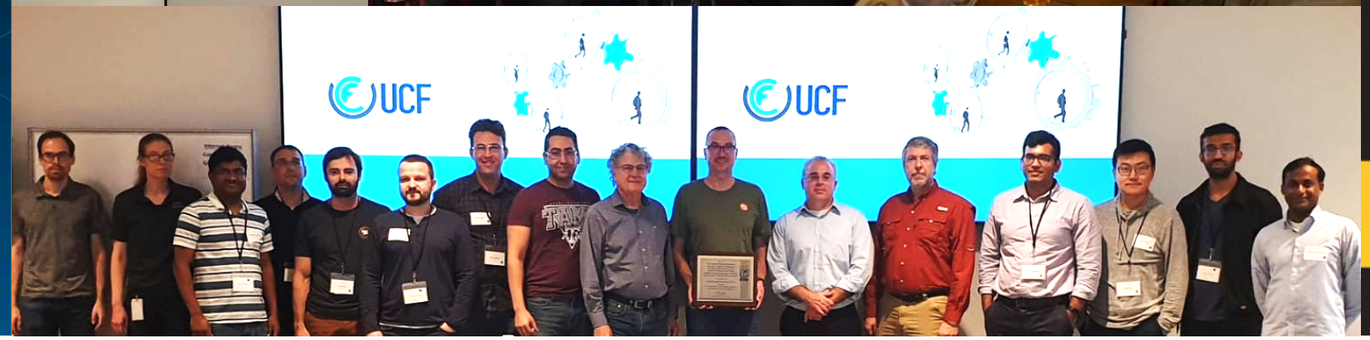

UCX High-level Overview

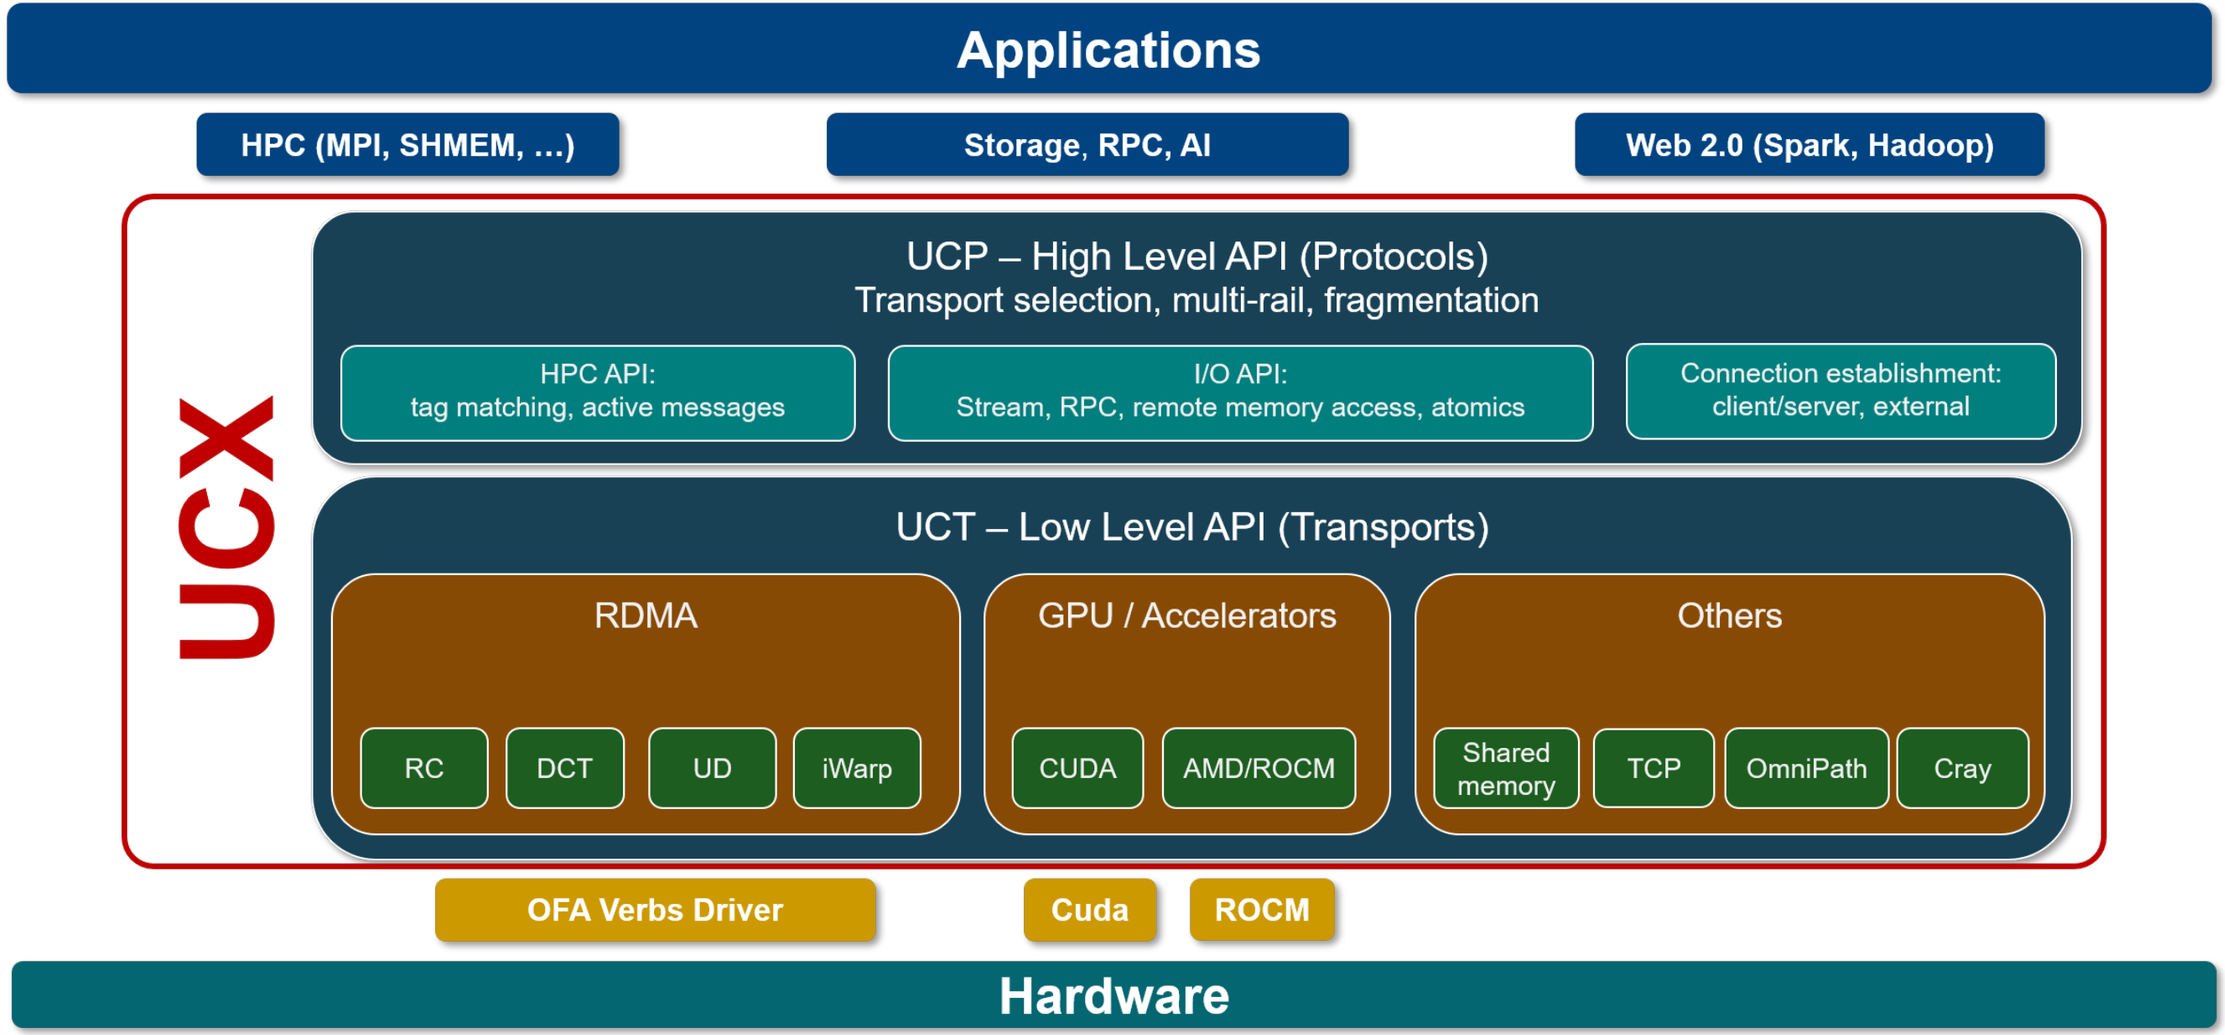

UCX Performance-portability

- Support for x86_64, Power 8/9, Arm v8
- U-arch tuned code for Xeon, AMD Rome/Naples, Arm v8 (Cortex-A/N1/ThunderX2/Huawei, Fujitsu A64FX)
- First class support for AMD and Nvidia GPUs
- Runs on Servers, Raspberry PI like platforms, SmartNIC, Nvidia Jetson platforms, etc.

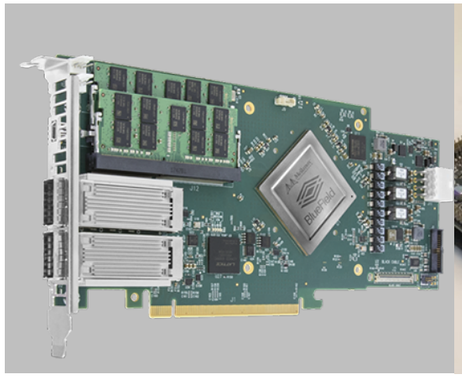

BlueField SmartNIC

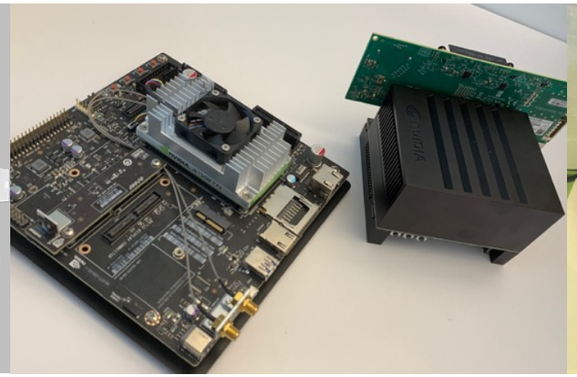

NVIDIA Jetson

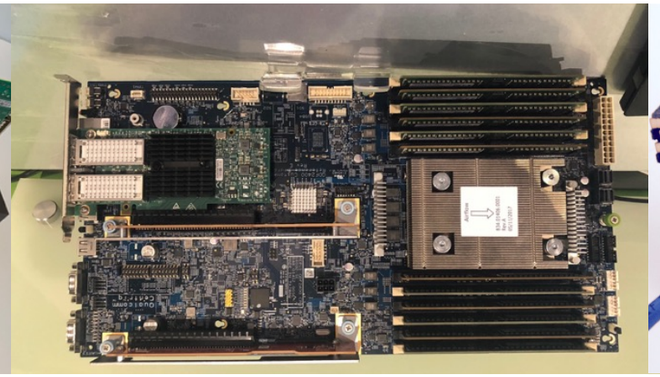

Arm ThunderX2

Odroid C2

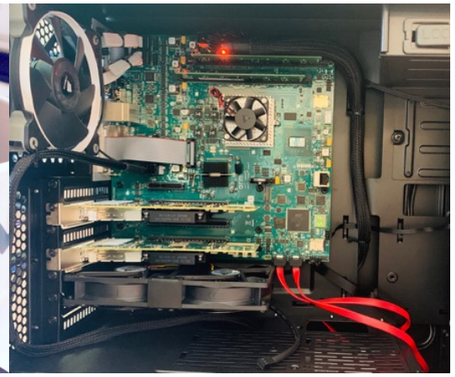

N1 SDP

We Love Testing

Over 100,000 tests per commit
220,000 CPU hours per release

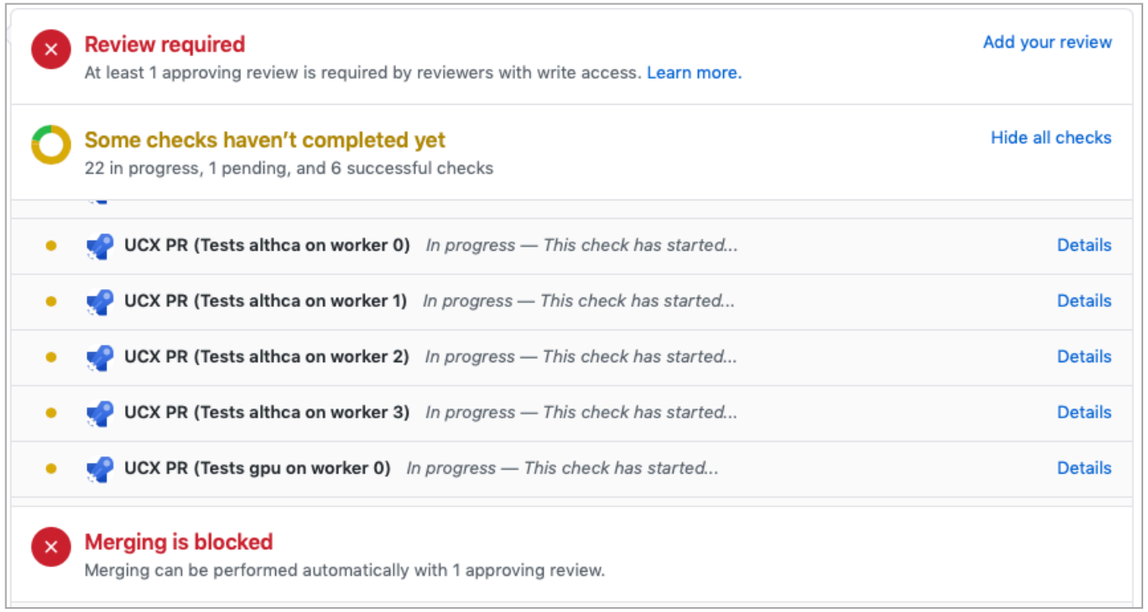

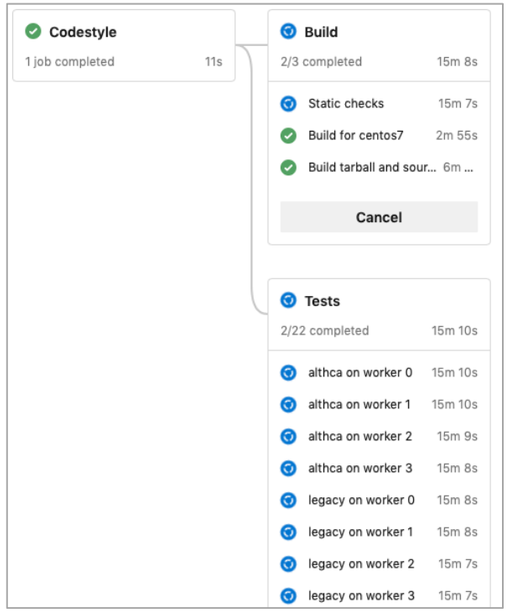

Annual Release Schedule

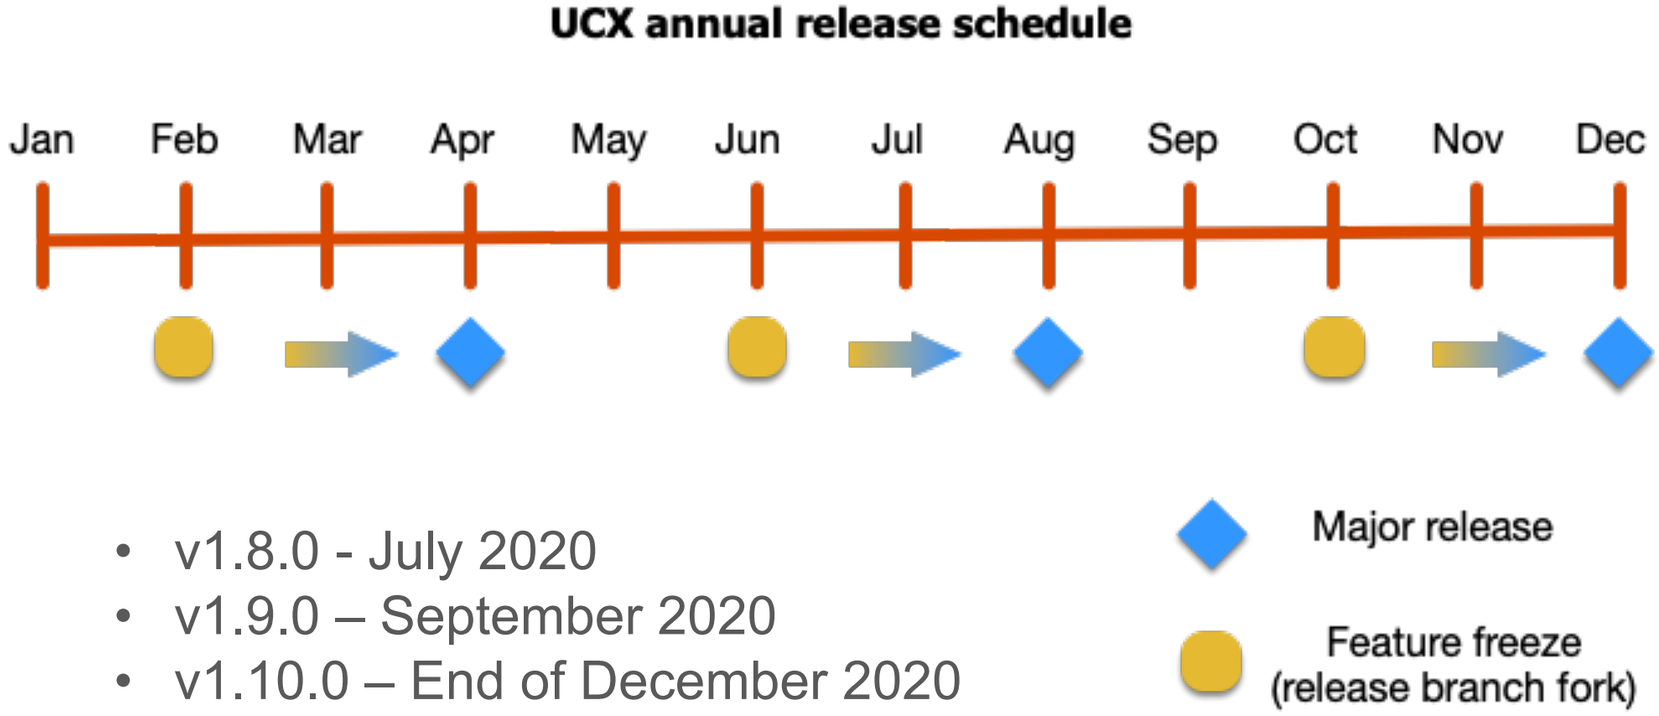

- v1.8.0 - July 2020
- v1.9.0 – September 2020
- v1.10.0 – End of December 2020

UCX Latest Releases

- **1.9.0 (September 19, 2020) - https://github.com/openucx/ucx/releases/tag/v1.9.0**
- **UCX Core**
  - Added a **new class of communication APIs** '*_nbx' that enable API extendability while preserving ABI backward compatibility
  - Added **asynchronous event** support to UCT/IB/DEVX
  - Added support for **latest NVIDIA CUDA** library version
  - Added support for **AMD ROCm 3.7 and above**
  - Added **new tests** for AMD ROCm
  - Added performance optimization for **Fujitsu A64FX** with InfiniBand
  - Added support for **relaxed-order PCIe access** in IB RDMA transports
  - Added new **TCP connection manager**
  - Added **flow control** for RDMA read operations
  - Improved **performance in active message** flow
- **UCX Java (API Preview)**
  - Added support for UCX shared library loading from both classpath and LD_LIBRARY_PATH
  - Added configuration map to ucp_params to be able to set UCX properties programmatically

UCX Users

- MPI implementations: MPICH, Open MPI, Mellanox MPI, Huawei MPI
- PGAS: GasNET
- OpenSHMEM: OSSS SHMEM, Sandia SHMEM, Open MPI SHMEM
- Charm++
- RAPIDS / DASK
- NVIDIA's NCCL

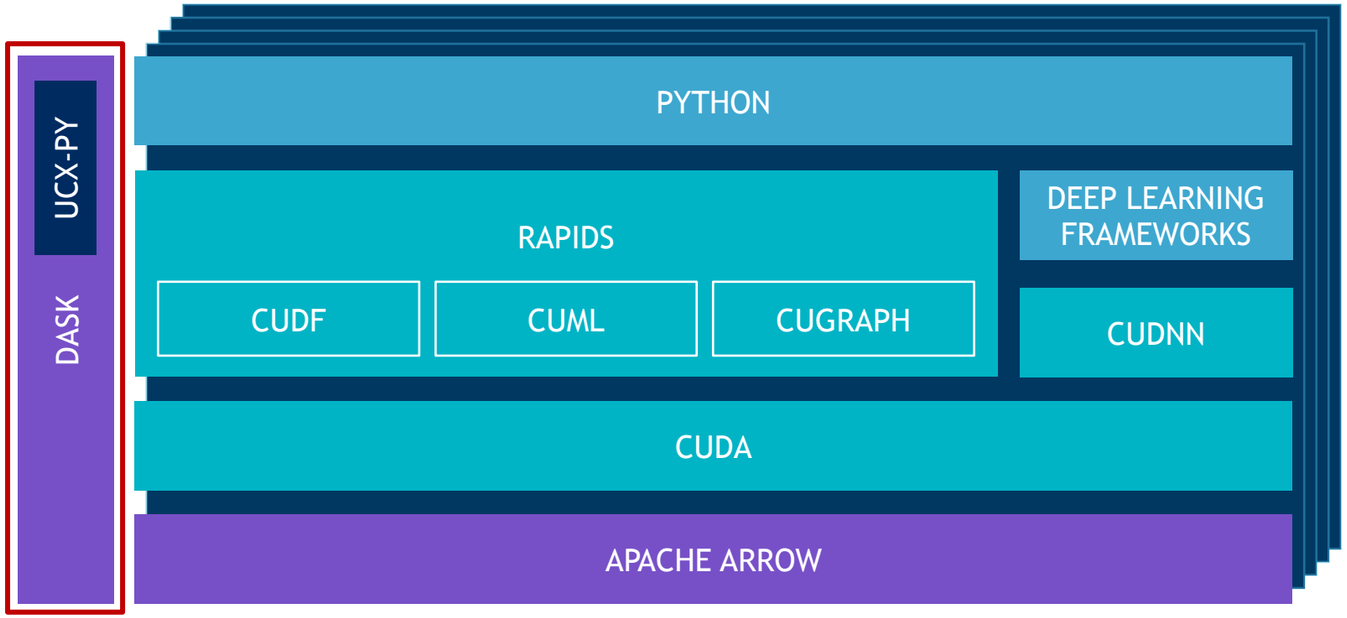

Diagram courtesy of Nvidia

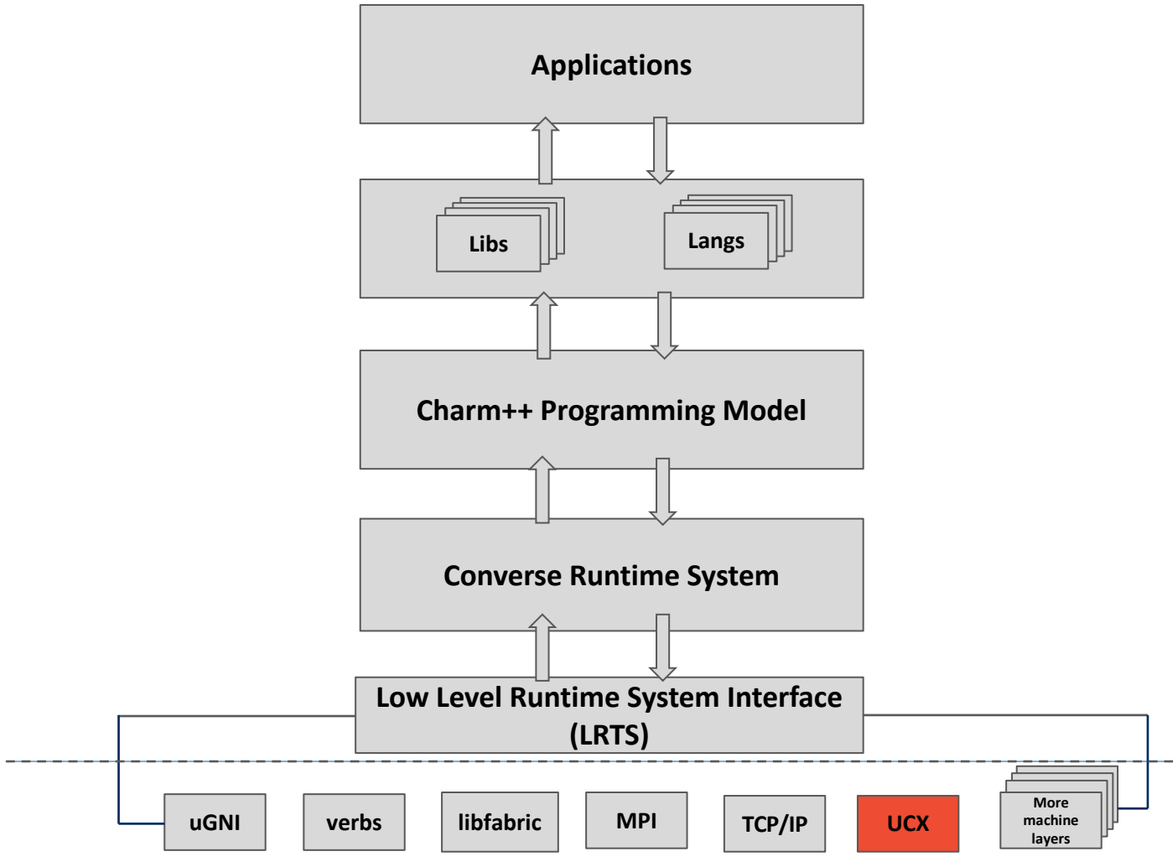

Diagram courtesy of Nitin Bhat @ Charmworks Inc

UCX – Useful links

- Code
  - https://github.com/openucx/
- Website
  - www.openucx.com
- Mailing list
  - https://elist.ornl.gov/mailman/listinfo/ucx-group
- Contributor agreement
  - https://www.openucx.org/license/
- User documentation
  - https://openucx.readthedocs.io/

UCF Project Incubator

Collective Communition API
Web https://www.ucfconsortium.org/projects/ucc/
GIT https://github.com/openucx/ucc
https://github.com/openucx/xucg
https://github.com/openucx/xccl
https://github.com/openucx/ucc_spec
https://github.com/openucx/torch-ucc

API for Smart NIC & DPU programmability
Web https://www.ucfconsortium.org/projects/opensnapi/
GIT https://github.com/openucx/shmem-opensnapi

info@ucfconsortium.org

UCF Project Incubator

UCX for Apache Spark™

A high-performance, scalable and efficient ShuffleManager plugin for Apache Spark™
Web https://www.ucfconsortium.org/projects/sparkucx/
GIT https://github.com/openucx/shmem-opensnapi

An effort to create a new metric for ranking HPC and AI system performance and capabilities
Web https://www.ucfconsortium.org/projects/hpca-benchmark/

UCD

Advanced date serialization/deserialization engine

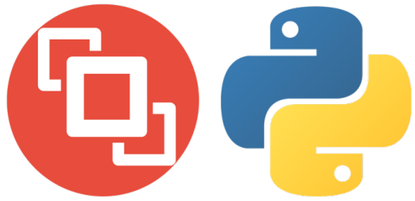

UCX-Py - Python bindings for UCX
GIT https://github.com/rapidsai/ucx-py

info@ucfconsortium.org

UCF

Unified
Communication
Framework

Thank You