codeplay

Enable AI & HPC to be Open, Safe and Accessible to All

Bringing SYCL™ to Ampere Architecture

Steffen Larsen – Staff Software Engineer

Partners

BERKELEY LAB

Argonne NATIONAL LABORATORY

IWOCL 2021 – Poster session

The Perlmutter Supercomputer

- Upcoming pre-exascale supercomputer at Lawrence Berkeley National Laboratory (Berkeley Lab.)
- Named after Saul Perlmutter, astrophysicist at Berkeley Lab and 2011 Nobel Laureate.
- HPE Cray system with CPU-only and GPU-accelerated nodes.
- A total of 6000 NVIDIA® A100 GPUs.

SYCL at (Pre-)Exascale

- Argonne National Laboratory (ANL) employs NVIDIA A100 GPU nodes in their ThetaGPU system.
- ANL's upcoming Aurora exascale supercomputer will support SYCL™ on Intel® hardware.
- SYCL on Perlmutter and ThetaGPU would mean portability and synergy with Aurora.

DPC++ for CUDA®

- Partnership between Codeplay, Berkeley Lab, and ANL.
- Goals:
  - Expand DPC++ for CUDA support for SYCL 2020 features.
  - Expose NVIDIA Ampere features in DPC++ for CUDA.
  - Optimize DPC++ for CUDA for NVIDIA Ampere hardware.

SYCL 2020 and Beyond

- Ensuring SYCL 2020 feature support in DPC++ for CUDA, including:
  - Unified Shared Memory (USM.)
  - Reductions.
  - Group and sub-group operations.
- Planned SYCL extensions:
  - Asynchronous barriers.
  - Tensor Core support.
- Better DPC++ for CUDA multi-device support.

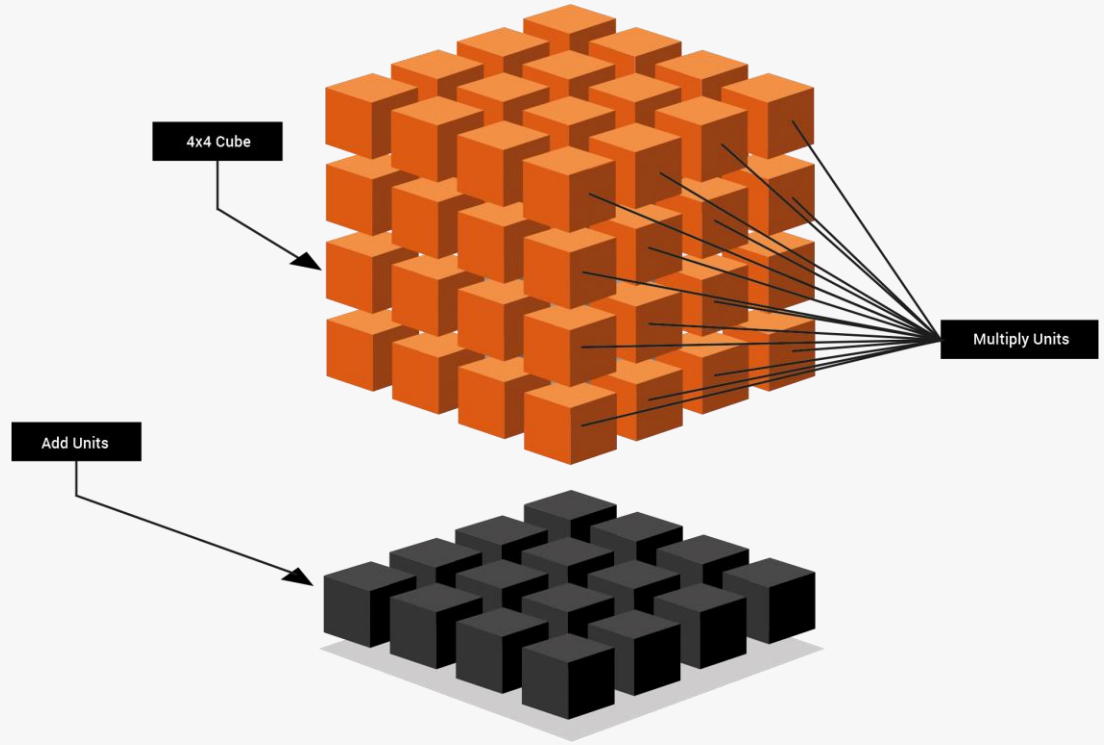

We're Hiring!

codeplay.com/careers/

codeplay

Enable AI & HPC to be Open, Safe and Accessible to All

@codeplaysoft

info@codeplay.com

codeplay.com