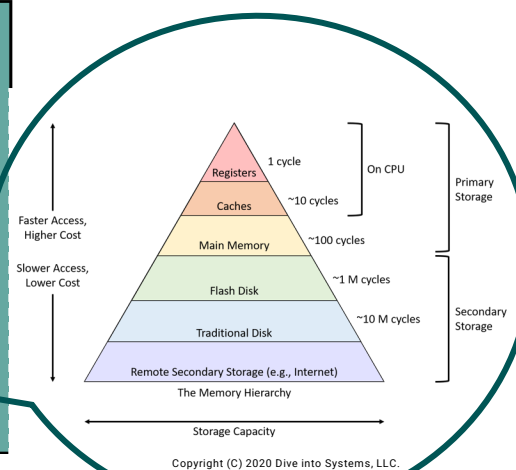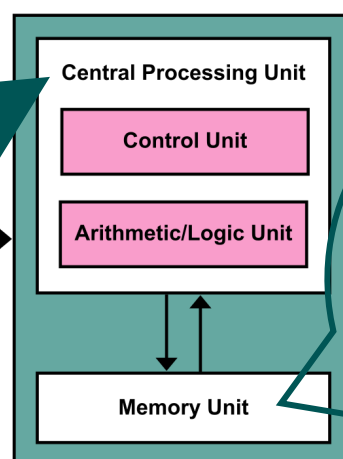MPCDF

MAX PLANCK
COMPUTING & DATA FACILITY

**VORBESPRECHUNG
SEMINAR
(PERFORMANCE) PORTABLE PROGRAMMING OF HPC
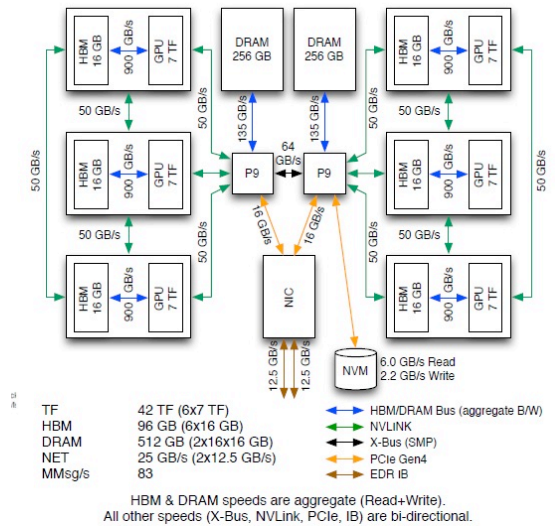APPLICATIONS (IN0014, IN2107)**

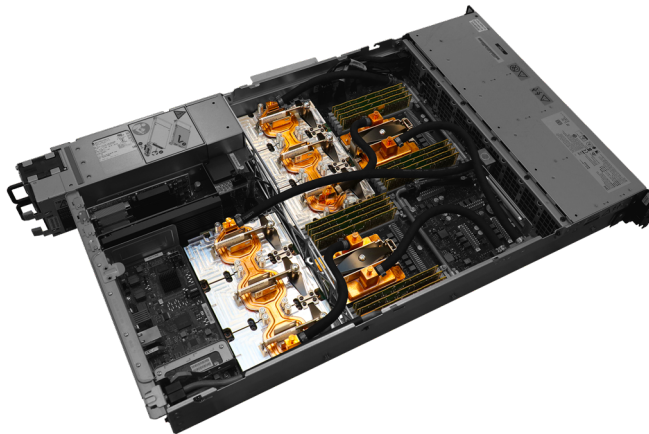**Prof. Dr. Erwin Laure**

1

FROM VON NEUMAN ARCHITECTURES TO MODERN HPC SYSTEMS

- **Multi-Core**
  - **E.g. 36-core Intel IceLake**
- **Lots of Optimizations**
  - **Pre-fetch**
  - **Branch prediction**
  - **FMA**
  - **Vector**
  - **Etc.**
- **Other features**
  - **Encryption**
  - **Viz**
  - **Etc.**

Central Processing Unit

Control Unit

Arithmetic/Logic Unit

Memory Unit

Faster Access, Higher Cost

Slower Access, Lower Cost

Registers — 1 cycle
Caches — ~10 cycles
Main Memory — ~100 cycles
Flash Disk — ~1 M cycles
Traditional Disk — ~10 M cycles
Remote Secondary Storage (e.g., Internet)

On CPU

Primary Storage

Secondary Storage

The Memory Hierarchy

Storage Capacity

Copyright (C) 2020 Dive into Systems, LLC.

2

AND THEN WE ALSO ADD ACCELERATORS (GPUS)

3

AND USE MANY, REALLY MANY OF THESE NODES

- **Summit Supercomputer @ ORNL:**
  - 4.608 nodes
  - 200 PF performance
  - 13 MW power consumption

4

(SOME) CHALLENGES IN PROGRAMMING THESE SYSTEMS

- **Level of parallelism**
  - $O(10^9)$ FPUs

- **Hardware heterogeneity**
  - CPUs, GPUs, other
  - HBM, SSD, object store

- **Programming/Performance Portability**

- **Novel numerical/methodological approaches**

5

THE GOOD OLD TIMES

- **Programms written in C/C++ or Fortran**

- **MPI (Message Passing Interface) for moving data across distributed memory**

- **OpenMP for expressing parallelism on shared memory**

6

Programming Landscape Today

CRAY
EPiGRAM HS

CUDA   OpenMP   ROCm/HIP

HSA   Programmer APIs   OpenCL

MPI   SYCL/One API

memkind   OpenACC   PCI   NVMe   Accelerators

Programming Frameworks   Memory APIs   Bus/Interconnect

tmpfs   Low-level interfaces   POSIX shared memory

OS   persistence   XPMEM

NUMA
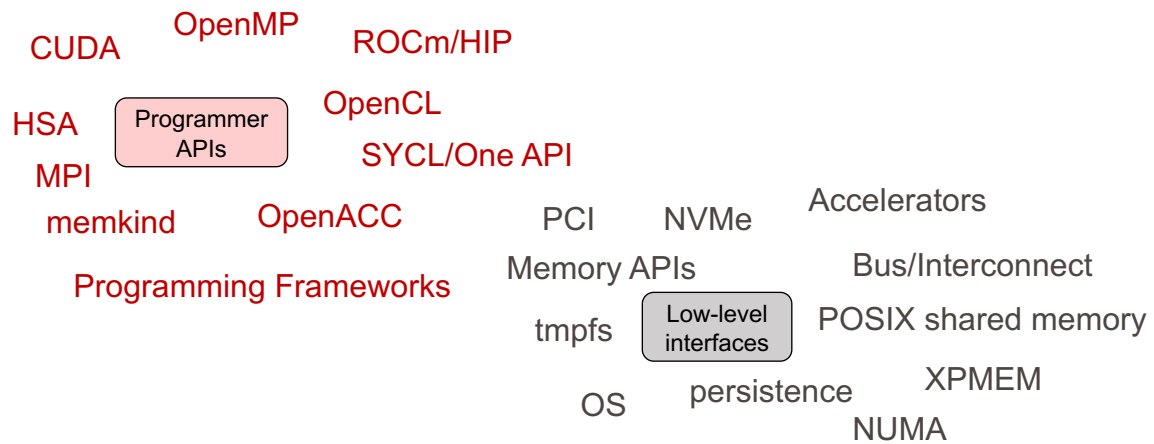
© 2019 Cray Inc.

7

GOALS OF THE SEMINAR

- **Investigate different programming approaches particularly wrt portability and performance portability**
  - Focus on heterogeneous architectures (GPUs, shared/distributed memory)

- **Examples:**
  - CUDA, HIP, OpenMP, Sycl
  - Performance portability frameworks (Kokkos, Raja, …)
  - Frameworks for AI
  - Abstractions & tools for memory/IO-management

8

SEMINAR ORGANIZATION

- **Kick-off meeting late April**
  - Final definition and selection of topics

- **Seminar paper (6-8 pages)**
  - Literature study (scientific papers! Min 3-4)
  - Main concepts (pros & cons) plus (where possible) experiences from real applications
  - Peer reviewed by seminar participants
    - 2nd week of June

- **Presentation (~15 mins)**
  - Workshop in late June/early July

- **Tutors will help in case of questions/problems**
  - Review paper/presentation drafts

- **Grading**
  - 40% paper, 40% presentation, 20% review
  - All needs to be positive

- **Prerequisits**
  - Understanding of parallel programming (e.g. IN2147)

MAX PLANCK COMPUTING AND DATA FACILITY | NAME

TITLE | DD.MM.YYYY   9

9