MAX PLANCK
COMPUTING & DATA FACILITY

MPCDF

VORBESPRECHUNG
SEMINAR ON EFFICIENT PROGRAMMING OF HPC SYSTEMS - FRAMEWORKS AND ALGORITHMS (IN2107)

Prof. Dr. Erwin Laure

FROM SIMPLE VON NEUMANN 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.

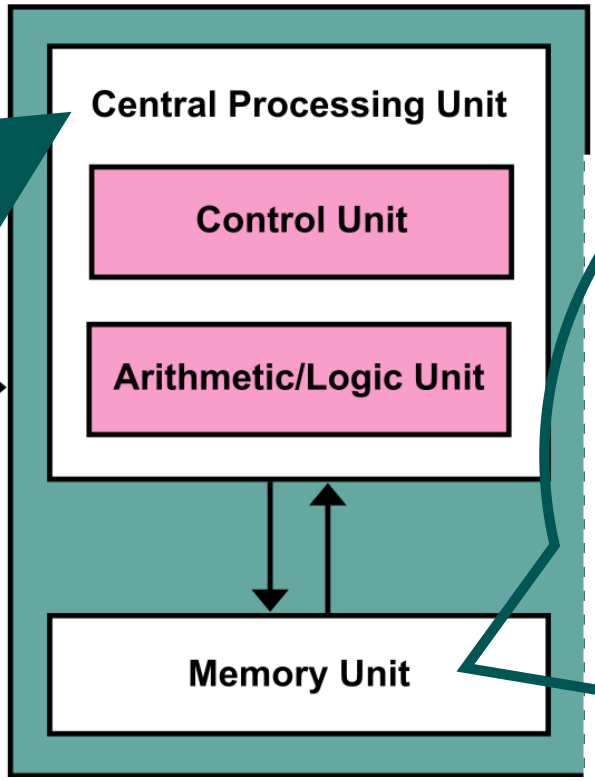

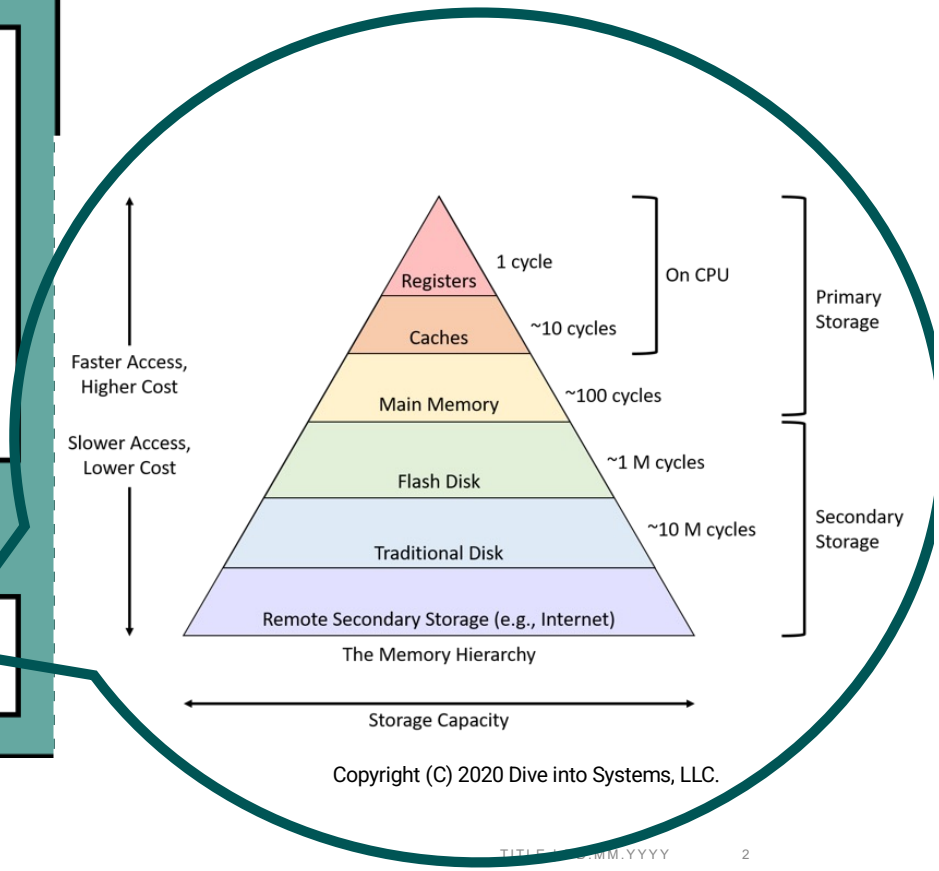

AND THEN WE ALSO ADD ACCELERATORS (GPUS)

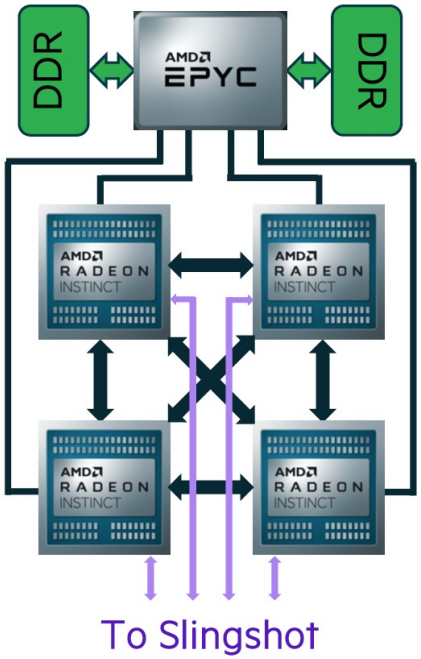

AMD GPU (ORNL)

2 nodes per blade

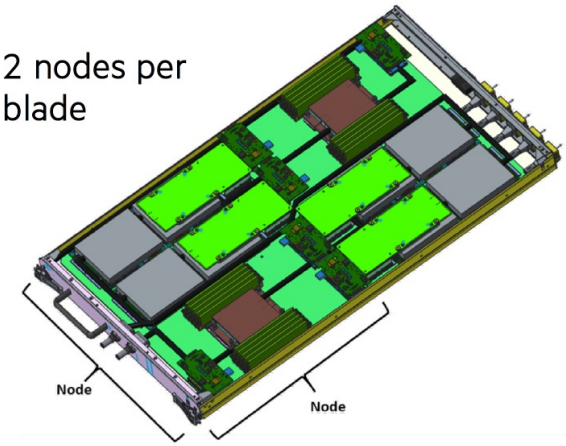

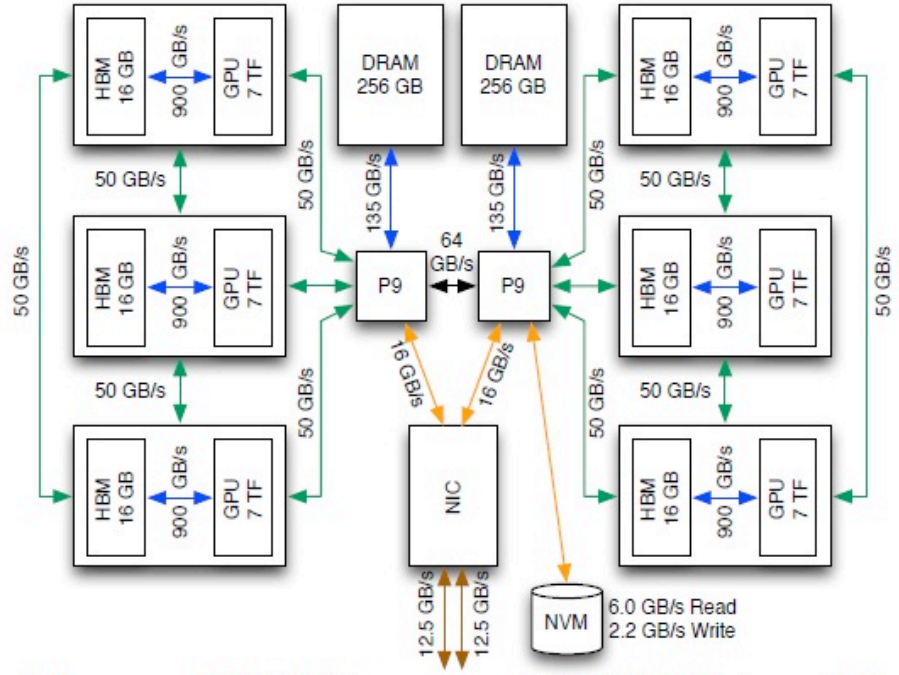

TF	42 TF (6x7 TF)
HBM	96 GB (6x16 GB)
DRAM	512 GB (2x16x16 GB)
NET	25 GB/s (2x12.5 GB/s)
MMsg/s	83

HBM/DRAM Bus (aggregate B/W)
NVLINK
X-Bus (SMP)
PCIe Gen4
EDR IB

HBM & DRAM speeds are aggregate (Read+Write).
All other speeds (X-Bus, NVLink, PCIe, IB) are bi-directional.

AND USE MANY, REALLY MANY OF THESE NODES

- **Frontier Supercomputer @ ORNL:**
  - 9.472 nodes
  - 1,1 EF performance
  - 21 MW power consumption
  - in total over 9 M cores (mostly GPU)

(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**

THE GOOD OLD TIMES

- Programms written in Fortran (or C/C++)
- MPI (Message Passing Interface) for moving data across distributed memory
- OpenMP for expressing parallelism on shared memory

Programming Landscape Today

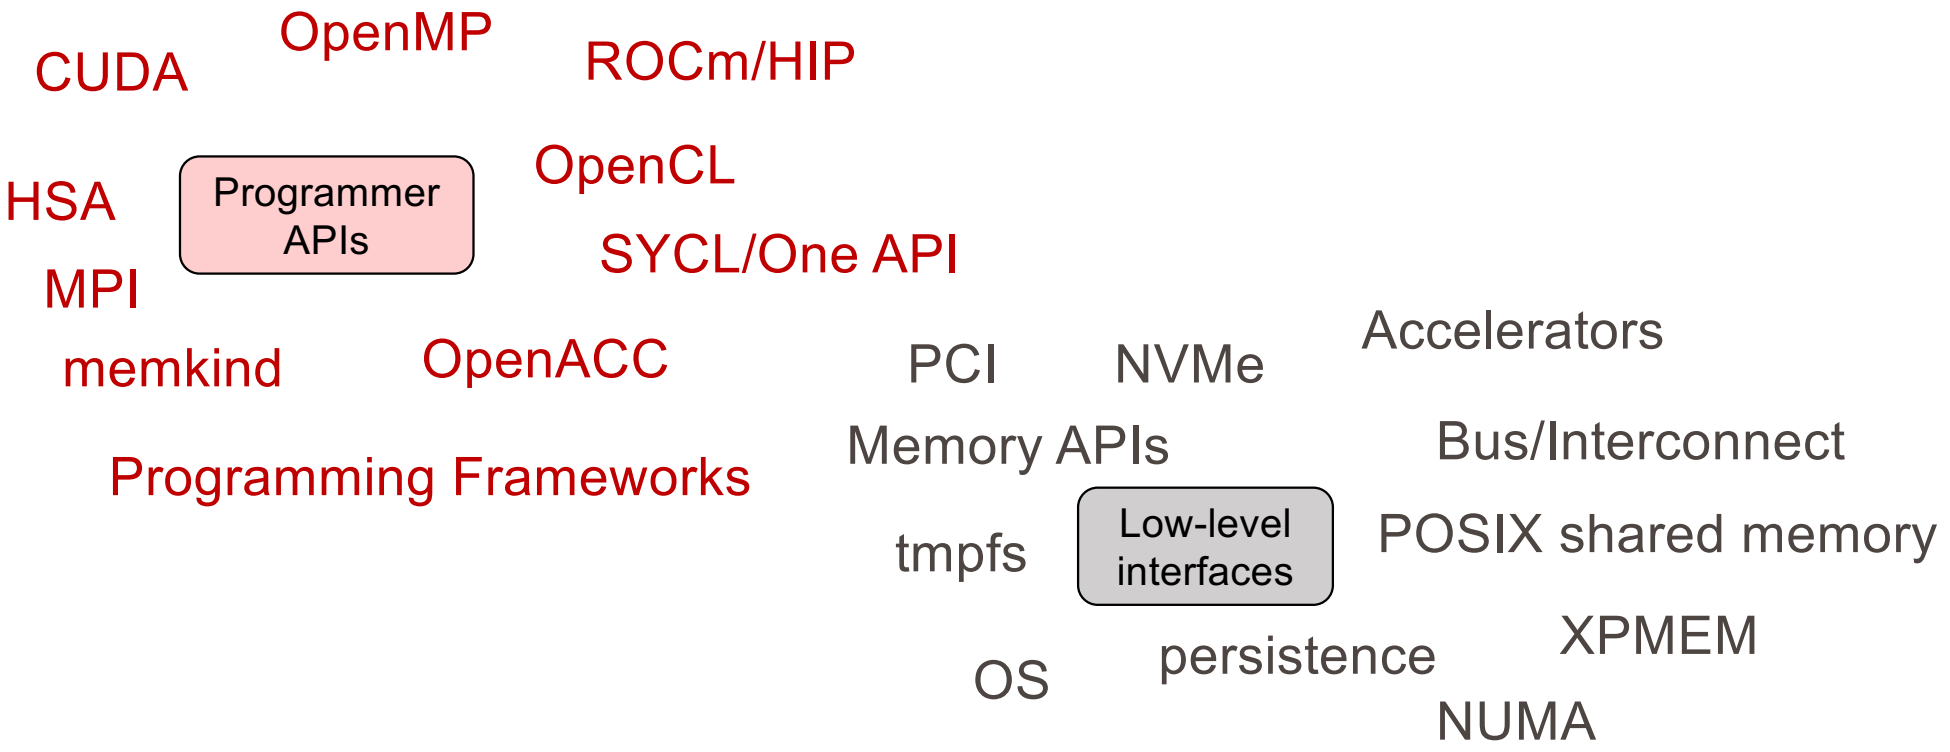

GOALS OF THE SEMINAR

- **Investigate techniques, frameworks, algorithms to efficiently program such systems**
  - Focus on heterogeneous architectures (GPUs, shared/distributed memory)

- **Examples:**
  - High-level frameworks (Kokkos, Alpaka, Cabana, PETSc, etc.)
  - Numerical libraries (SLATE, Gingko, heFFTe, etc.)
  - Mixed-precision and use of non-IEEE data formats
  - Data structures and layouts (AoS, SoA, AoSoA)
  - Adaptive Mesh Refinement (AMReX, p4est, etc.)
  - Adaptive (task) Parallelism (HPX, StarPU, Charm++, etc.)
  - Frameworks for AI (pytorch, tensorflow, etc.)

SEMINAR ORGANIZATION

- **Kick-off meeting 18. April (15-17)**
  - 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 July 5th (whole day; if needed also on July 4th
- **Tutors will help in case of questions/problems**
  - Provide help at all stages
  - Review paper/presentation drafts
  - Mandatory to discuss concepts with them
- **Grading**
  - 40% paper, 40% presentation, 20% review
  - All needs to be positive
- **Prerequisits**
  - Understanding of parallel programming (e.g. IN2147)