SC25 PANEL SESSIONS

Hardware Modularity for Practical Heterogeneous HPC

System–Level Modularity

Hans-Christian Hoppe, Forschungszentrum Jülich

SC25 ST. LOUIS hpc ignites.

Practical Heterogeneity in HPC/AI – Opportunities and Challenges

JÜLICH Forschungszentrum

Facts of life from a HPC/AI center PoV

- Wide variety of workloads with diverging characteristics
- Significant amount of “legacy” codes PLUS rapidly evolving “new codes on the block”
- Need to drive down capital expense and energy consumption
- Systems must be good for > 5 years, better “chose wisely”

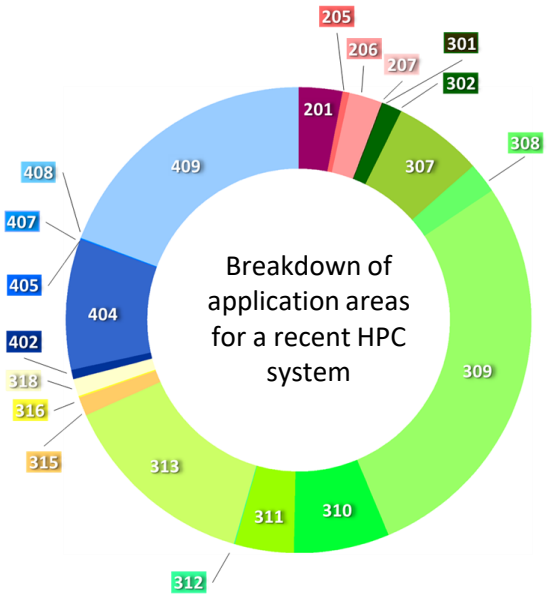

Increasing amount of architecture choices

- Heterogeneity enables “best of breed” combinations matching important codes/domains
- The plurality of workloads must be covered, avoiding
  - Too much “dark silicon”
  - Inefficient execution
  - Unaffordable code modifications / recoding

O(100) cores

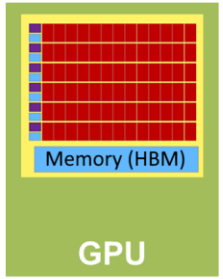

O(1000) compute units

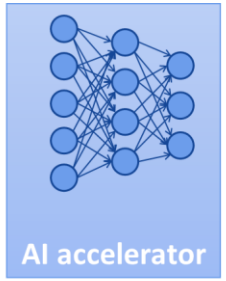

Deep Learning engines

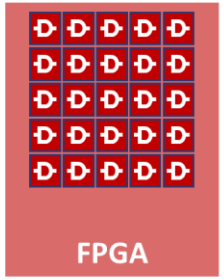

O(1000000) gates

O(10) gates

Modularity at the HPC/AI System Level

- Compose modular heterogenous system from (large-scale) component system
- Network federation facilitates communication and coordination
- Modules can be added/extended/replaced

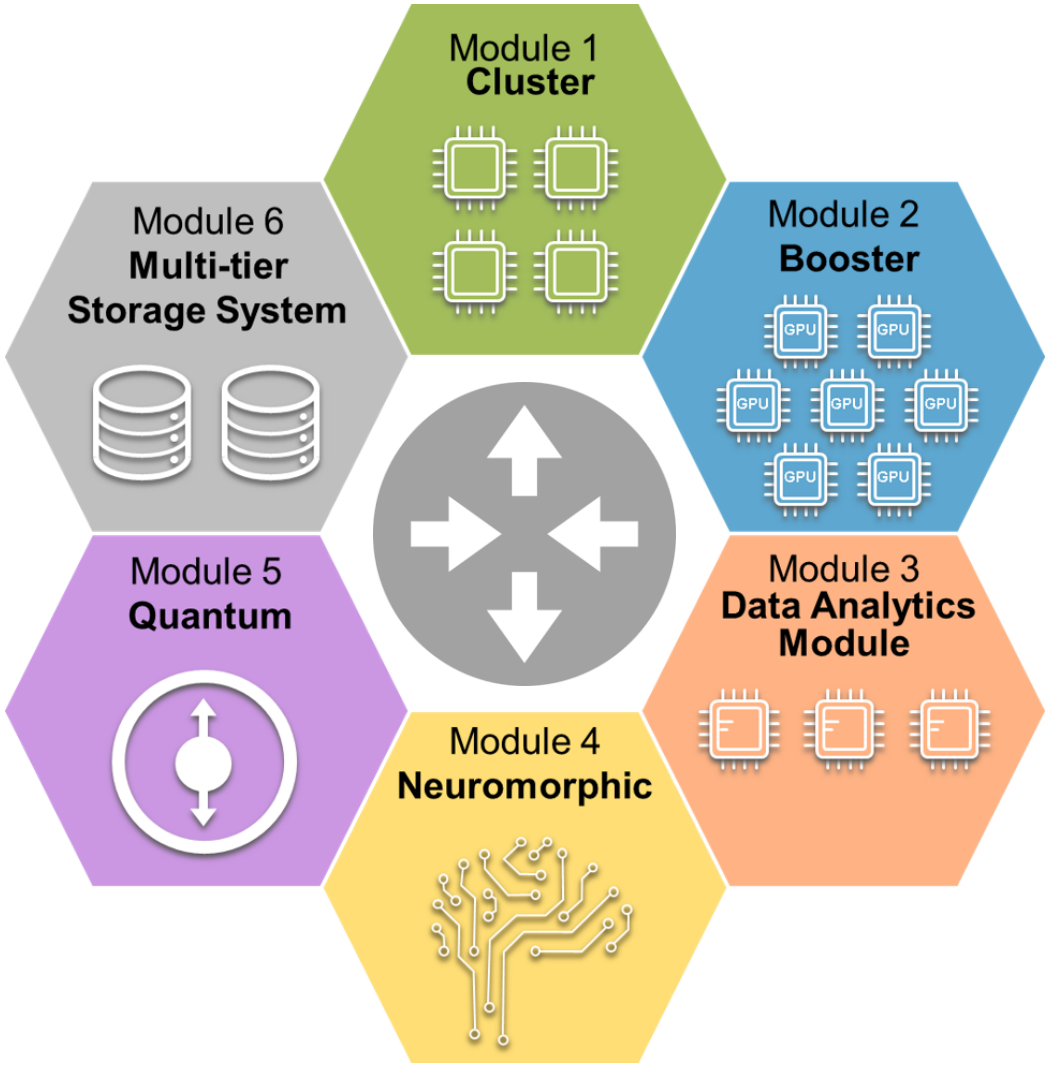

- **Suarez** et al. *"Modular Supercomputing Architecture – A Success Story of European R&D",* ETP4HPC White Paper. (2022) Available at https://www.etp4hpc.eu/white-papers.html#msa.
- **Suarez** et al., "*Modular Supercomputing Architecture: from idea to production*", Chapter 9 in Contemporary High Performance Computing: from Petascale toward Exascale, Volume 3, p 223-251, CRC Press. (2019)

Workloads on a Modular System

Just on one module

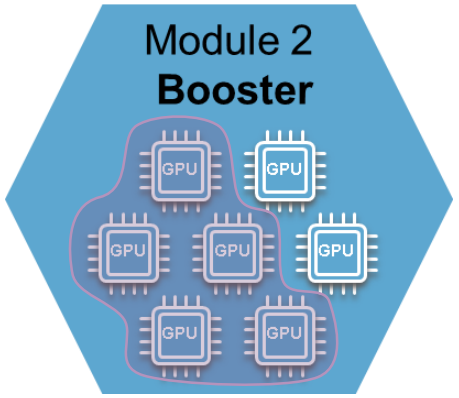

Typical of **tightly-coupled**, high-scaling **applications** (e.g. dense M×M)

Workflows

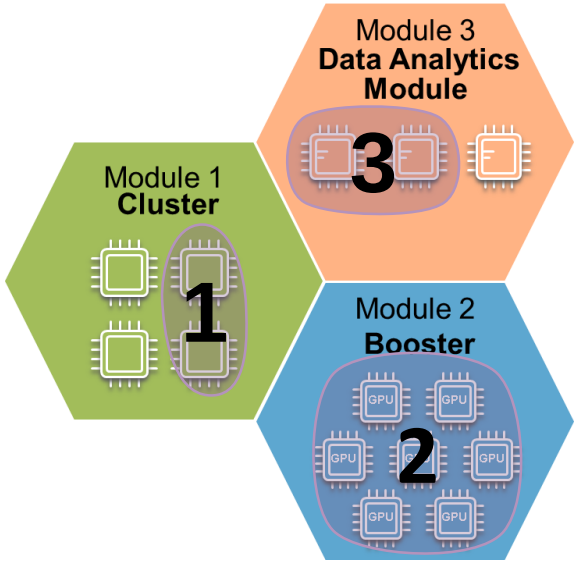

Typical of complex **application workflows** (MPMD applications, e.g. pre-processing→ simulation→data analysis)

Across modules

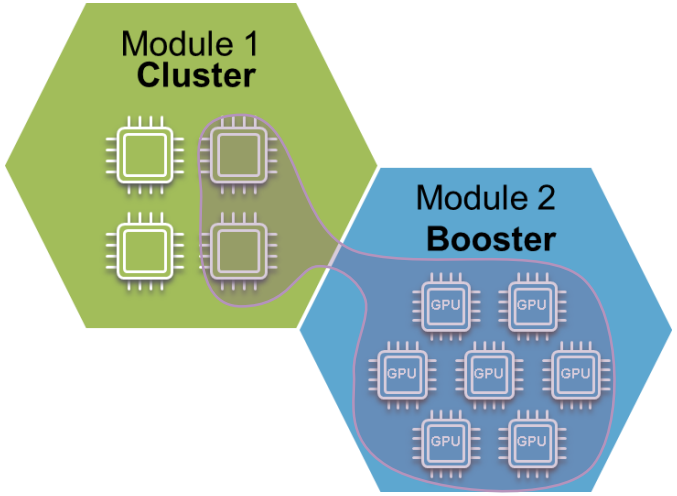

Typical for **multi-physics** or multi-scale applications (e.g. coupled climate models using MPI)

Modular System Evolution