UMBC High Performance Computing Facility (HPCF)

Complete information at the HPCF webpage www.umbc.edu/hpcf:

- Research including public relations: list of projects, list of publications, posters
- Resources for users: usage policies, technical details, usage information including tutorials and sample code, link to consulting support, supporting material
- Point of contact: Matthias K. Gobbert, *gobbert@umbc.edu*, 410-455-2404

- **HPCF is the community-based, interdisciplinary core facility for scientific computing and research on parallel algorithms at UMBC.**

- The current 240-node cluster **maya** contains 156 nodes connected by a high-speed quad-data rate (QDR) InfiniBand network for research on parallel algorithms. 84 of these nodes feature two quad-core 2.6 GHz Intel Nehalem X5550 CPUs and 24 GB memory each, while the newest 72 nodes contain two eight-core 2.6 GHz Intel E5-2650v2 Ivy Bridge CPUs and 64 GB memory. With the addition of the newest 72 nodes in 2013, HPCF boasts for the first time hybrid nodes, with 19 of these featuring each two high-end NVIDIA K20 GPUs (graphics processing units) designed for double-precision arithmetic and another 19 nodes featuring each two cutting-edge 60-core Intel Phi 5110P accelerators. The remaining 84 nodes in maya are designed for fastest number crunching, with two quad-core 2.8 GHz Intel Nehalem X5560 CPUs and 24 GB memory, connected by a dual-data rate (DDR) InfiniBand network. All nodes are connected via InfiniBand to a central storage of more than 750 TB.

- The HPCF initiative originated in 2008 with an MRI proposal to the National Science Foundation (NSF) by over 20 faculty from more than 10 departments and research centers from all three colleges at UMBC. It built on the experience with a partnership between the Department of Mathematics and Statistics and the Division of Information Technology (DoIT) in jointly operating the 33-node cluster **kali** (purchased in 2003 with funds from an NSF SCREMS grant and UMBC) that had also been used by researchers from several other departments, notably in the College of Natural and Mathematical Sciences. The initial machine **hpc** in HPCF with 35 nodes was funded jointly by several participating faculty, the administration, and DoIT in 2008 and extended the partnership with DoIT to the entire UMBC community. Its complete replacement and significant expansion to the 86-node **tara** in 2009 was supported by NSF grants from the MRI and SCREMS programs, plus funding from individual researchers. The dramatic expansion to the 240-node **maya** in 2013 used a second NSF MRI grant, UMBC funds, individual researchers' funds, and a gift from NASA, the first-ever significant gift of computing equipment to UMBC.

- **Unique features of HPCF:**
  - Integration with education: Math 447 and 627 Introduction to Parallel Computing. *REU Site: Interdisciplinary Program in High Performance Computing* (www.umbc.edu/hpcreu) funded jointly by NSF and NSA as summer program for 24 supported students. Used as resource for other grants, e.g., NSF-funded Biology-Mathematics UBM@UMBC program (www.umbc.edu/ubm).
  - Individual user support available by 2 RAs funded by UMBC and via consulting through the *Center for Interdisciplinary Research and Consulting* (www.umbc.edu/circ) in the Department of Mathematics and Statistics.
  - HPCF Governance Committee comprised of members of the user community. Regular meetings of the user support team with DoIT staff on system management.
  - Scientific and parallel computing research visible on campus (more than 300 users total; 40 faculty research groups)
  - Interdisciplinary research opportunities among departments and with research centers (JCET, CUERE, IRC, etc.)

How does HPCF stack up?

- **Research-group (or department) owned computer:**
    - Size of system: 8 or 16 nodes
    - Network: basic Ethernet
    - Storage: few TB (not backed up, no redundancy)
    - System administration: by users (graduate student)
    - User support: none
- **College-based central computer:**
    - Size of system: 32 or 64 nodes
    - Network: varies
    - Storage: few TB (some redundancy, maybe backed up)
    - System administration: professional (part-time)
    - User support: varies, usually none
- **HPCF (after 2013 expansion):**
    - Size of system: 240 nodes, including 84 homogeneous
    - Network: high performance (low latency, wide bandwidth)
    - Storage: over 750 TB with redundancy, including 10 TB with backup
    - System administration: professional (redundant full-time)
    - User support: by local ticket and in-person locally available via consulting approach
- **National supercomputer centers (e.g., NSF funded EXCEDE members):**
    - Size of system: hundreds of nodes, but often only, e.g., 200 homogeneous
    - Network: high performance, at least in part of system
    - Storage: significant, at least 100s TB (redundancy, maybe mirrored, no full backup)
    - System administration: professional (several full-time)
    - User support: professional and/or student-worker remotely by e-mail or forum only
- **National labs (e.g., Department of Defense, Department of Energy):**
    - Size of system: into the thousands of nodes or specialized system
    - Network: high performance
    - Storage: very significant
    - System administration: professional (multiple full-time)
    - User support: irrelevant, since system not open to public