UMBC High Performance Computing Facility (HPCF)

UMBC

Contact: Matthias K. Gobbert, Department of Mathematics and Statistics, UMBC, gobbert@umbc.edu

HPCF Research Assistants: Carlos Barajas and Reetam Majumder, hpcf.umbc.edu

Governance Committee: Don Engel, Frank Ferraro, Daniel Lobo, Larrabee Strow, Jianwu Wang, Meilin Yu, Zhibo Zhang

HPCF Overview

The UMBC High Performance Computing Facility (HPCF) is the community-based, interdisciplinary core facility for scientific computing and research on parallel algorithms at UMBC. Started in 2008 by more than 20 researchers from ten academic departments and research centers from all academic colleges at UMBC, it is supported by faculty contributions, federal grants, and the UMBC administration.

System administration is provided by the UMBC Division of Information Technology, and users have access to consulting support provided by dedicated full-time graduate assistants. Researchers can contribute funding for long-term priority access.

HPCF Outcomes

Since HPCF's inception, over 400 users have profited from its computing clusters, including undergraduate and graduate students. The users generated over 300 publications, including 100 papers in peer-reviewed journals (including Nature, Science, and other top-tier journals in their fields), 50 refereed conference papers, and 50 theses. HPCF has enabled major funded initiatives including an REU Site and the UMBC CyberTraining initiative (cybertraining.umbc.edu).

Acknowledgments

NSF, NASA, NSA, UMBC, HPCF, CIRC

System: Heterogeneous Cluster taki with over 200 nodes

- CPU Cluster: **179 nodes totaling over 3,000 cores and 20 TB memory!**
  - 42 nodes with two 18-core Intel Skylake CPUs, 384 GB of memory, and EDR IB,
  - 49 nodes with two 8-core Intel Ivy Bridge CPUs, 64 GB of memory, and QDR IB,
  - 82 nodes with two 4-core Intel Nehalem CPUs, 24 GB of memory; and QDR IB;
- GPU Cluster:
  - 1 GPU node with four NVIDIA Tesla V100 GPUs connected by NVIDIA NVLink,
  - 18 CPU/GPU nodes with two CPUs and two NVIDIA K20 GPUs with QDR IB;
- Big Data Cluster: 8 nodes with two CPUs and 48 TB disk space with EDR IB.

All nodes are connected via InfiniBand to a central storage of more than 750 TB.

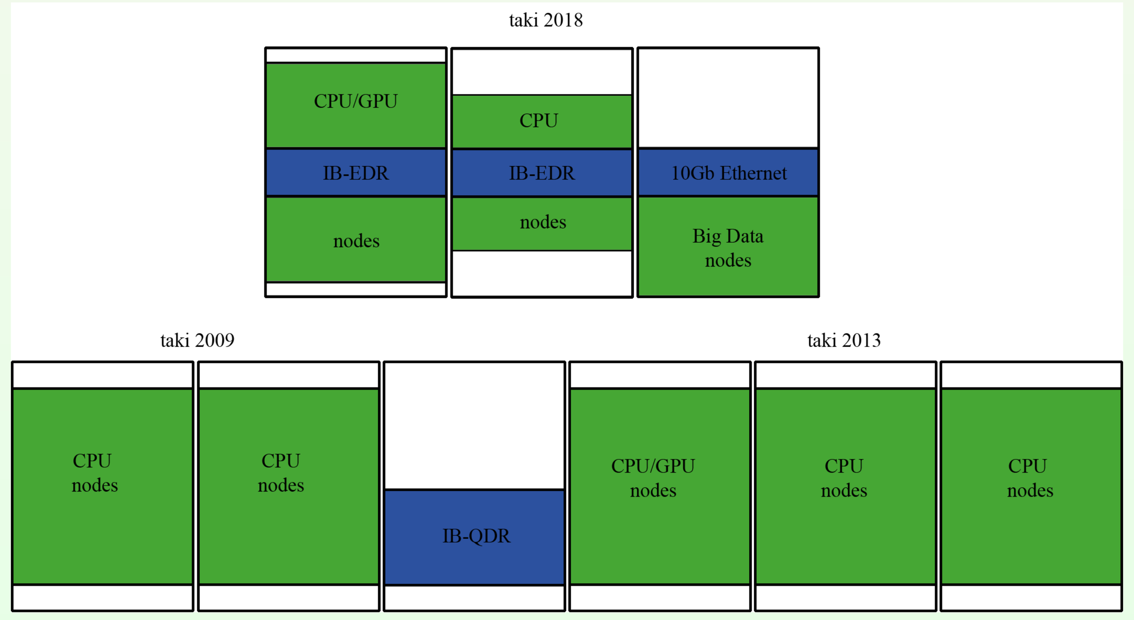

System layout

Big Data nodes

QDR IB front

2-D Poisson Equation Test Problem using MPI-Only Code on CPU Cluster

The Poisson equation in two spatial dimensions, discretized by finite differences and solved by CG method, on a $16{,}384 \times 16{,}384$ mesh with DOF over 250 million unknowns requires 26,316 iterations [HPCF–2019–1]. This test problem is memory-bound and indicative for behavior of many codes. Observed wall clock time in HH:MM:SS on 2018 CPU nodes by number of nodes and processes per node:

	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes
1 MPI process per node	15:28:12	07:41:17	03:47:46	01:54:36	00:56:34	00:27:46
2 MPI processes per node	07:39:45	03:52:01	01:54:28	00:57:21	00:28:05	00:13:37
4 MPI processes per node	03:56:37	01:58:27	00:58:52	00:29:23	00:14:29	00:07:04
8 MPI processes per node	02:03:22	01:01:43	00:30:56	00:15:18	00:07:37	00:03:47
16 MPI processes per node	01:22:16	00:41:19	00:20:39	00:10:18	00:05:07	00:02:30
32 MPI processes per node	01:10:50	00:35:34	00:17:53	00:09:03	00:04:33	00:02:18

- Speedup across nodes demonstrates quality of the high-performance interconnect.
- Speedup within shared-memory node demonstrates quality of multi-code CPUs.

User Support using CIRC

User support provided through a consulting approach in collaboration with the Center for Interdisciplinary Research and Consulting (circ.umbc.edu).

taki with Dr. Gobbert and HPCF RAs

CICR Problem on GPU Cluster

CICR simulation on $128 \times 128 \times 512$ mesh in 3-D with DOF 25,610,499 unknowns per time step [Huang, 2015]. Observed wall clock time of CUDA+MPI code against one 16-core CPU node:

nodes (GPU/node)	Time (speedup)
1 node (16 cores)	26:53:37
1 node (1 GPU)	15:32:34 (1.73)
1 node (2 GPUs)	08:18:18 (3.24)
2 nodes (16 cores)	13:56:38
2 nodes (1 GPU)	08:14:26 (3.26)
2 nodes (2 GPUs)	04:25:55 (6.07)
4 nodes (16 cores)	07:21:17
4 nodes (1 GPU)	04:20:56 (6.18)
4 nodes (2 GPUs)	02:28:00 (10.90)
8 nodes (16 cores)	03:54:47
8 nodes (1 GPU)	02:24:46 (11.15)
8 nodes (2 GPUs)	01:31:22 (17.66)
16 nodes (16 cores)	02:17:48
16 nodes (1 GPU)	01:30:06 (17.91)
16 nodes (2 GPUs)	01:06:17 (24.34)