

# UMBC High Performance Computing Facility



Contact: Matthias K. Gobbert, Department of Mathematics and Statistics, [www.umbc.edu/hpcf](http://www.umbc.edu/hpcf)

Research Assistants: Xuan Huang, Samuel Khuvis, Jonathan Graf, Jack Slettebak, Andrew Raim


Governance Committee: M. K. Gobbert, C. R. Menyuk, M. Olano, L. Sparling, L. Strow, I. Thorpe, C. Welty



## HPCF Introduction

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 three colleges, it is supported by faculty contributions, federal grants, and the UMBC administration. The facility is open to UMBC researchers at no charge. Researchers can contribute funding for long-term priority access. 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.


## Photos



HPCF2013 Front



QDR IB Front

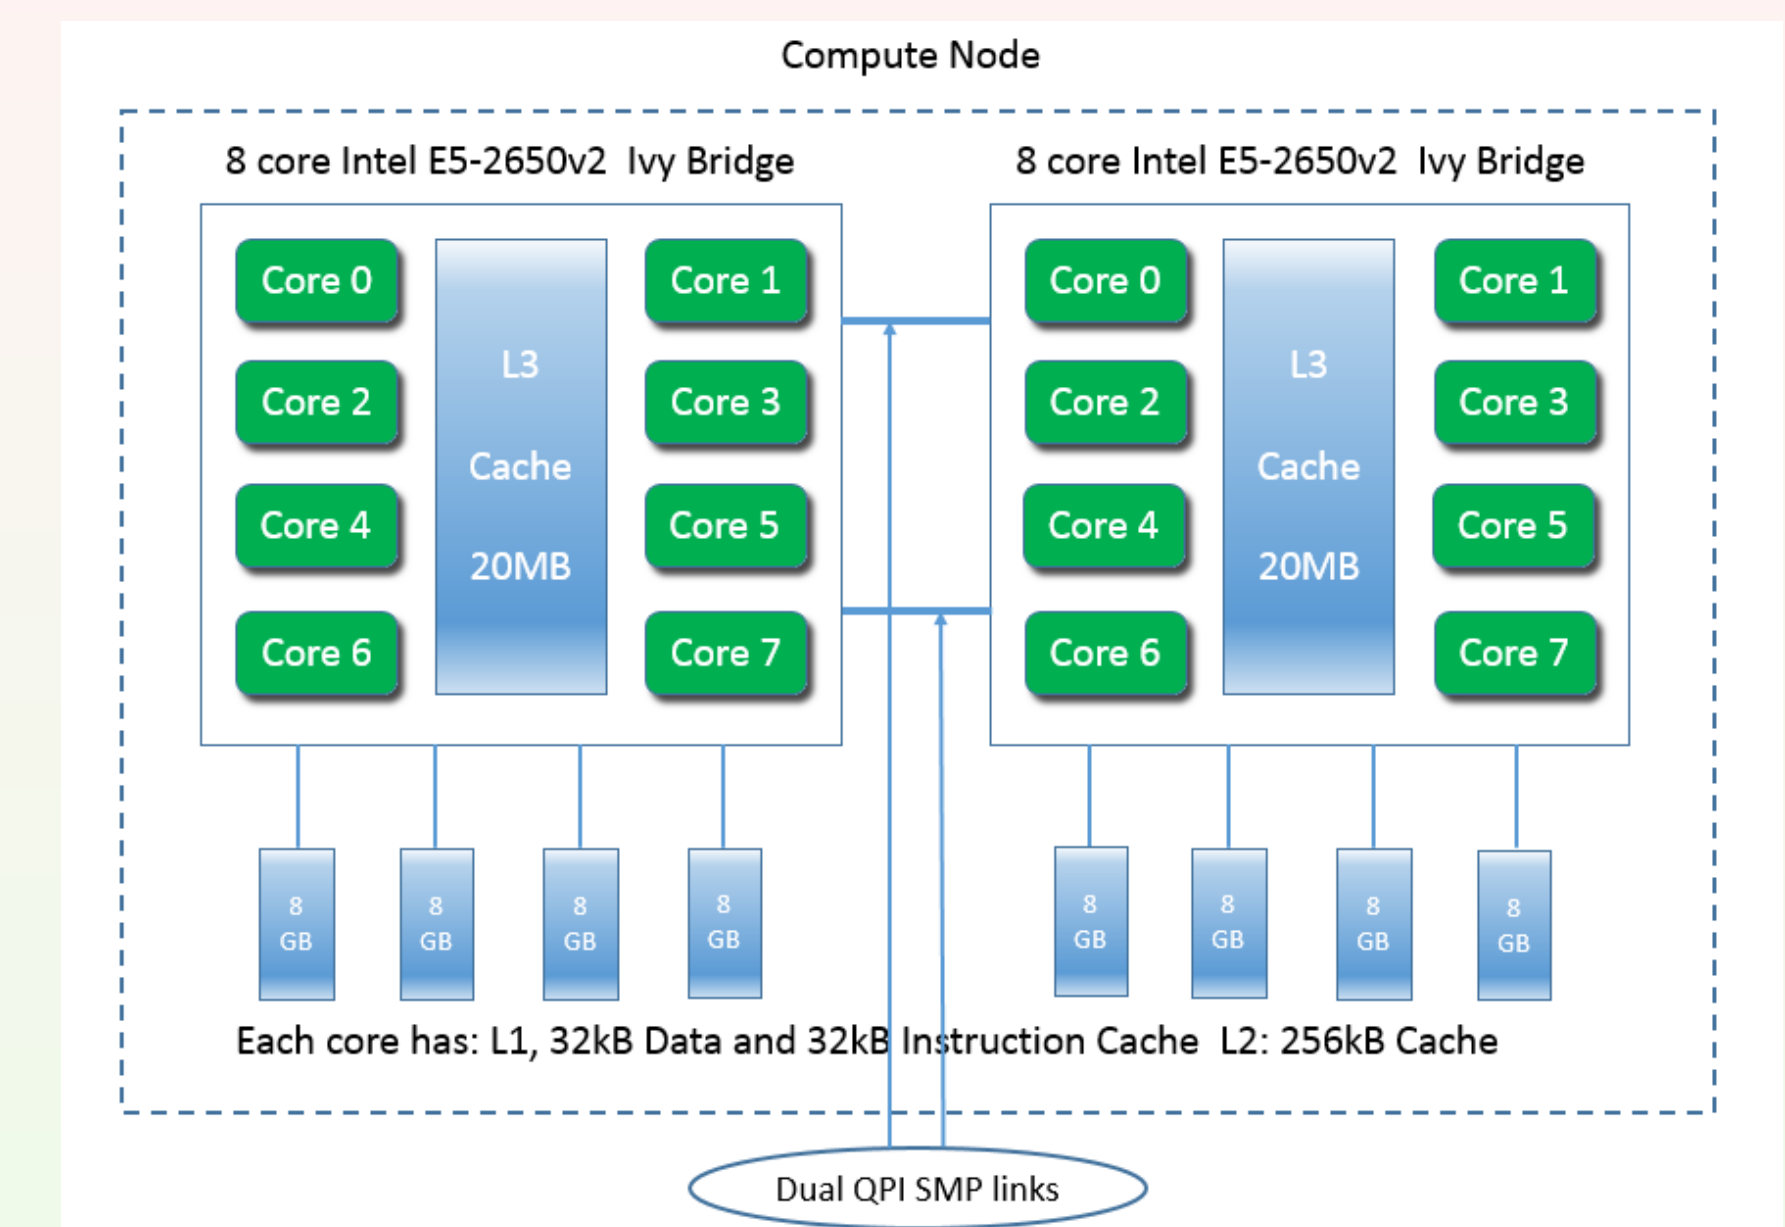


HPCF2013 Back



QDR IB Back

## HPCF2013: 2 Eight-Core CPUs



- Each node contains two eight-core Intel E5-2650v2 Ivy Bridge CPUs.
- The 64 GB of the node's memory are connected to each CPU via 4 memory channels.
- The two CPUs of a node are connected to each other by two QPI (quick path interconnect) links.

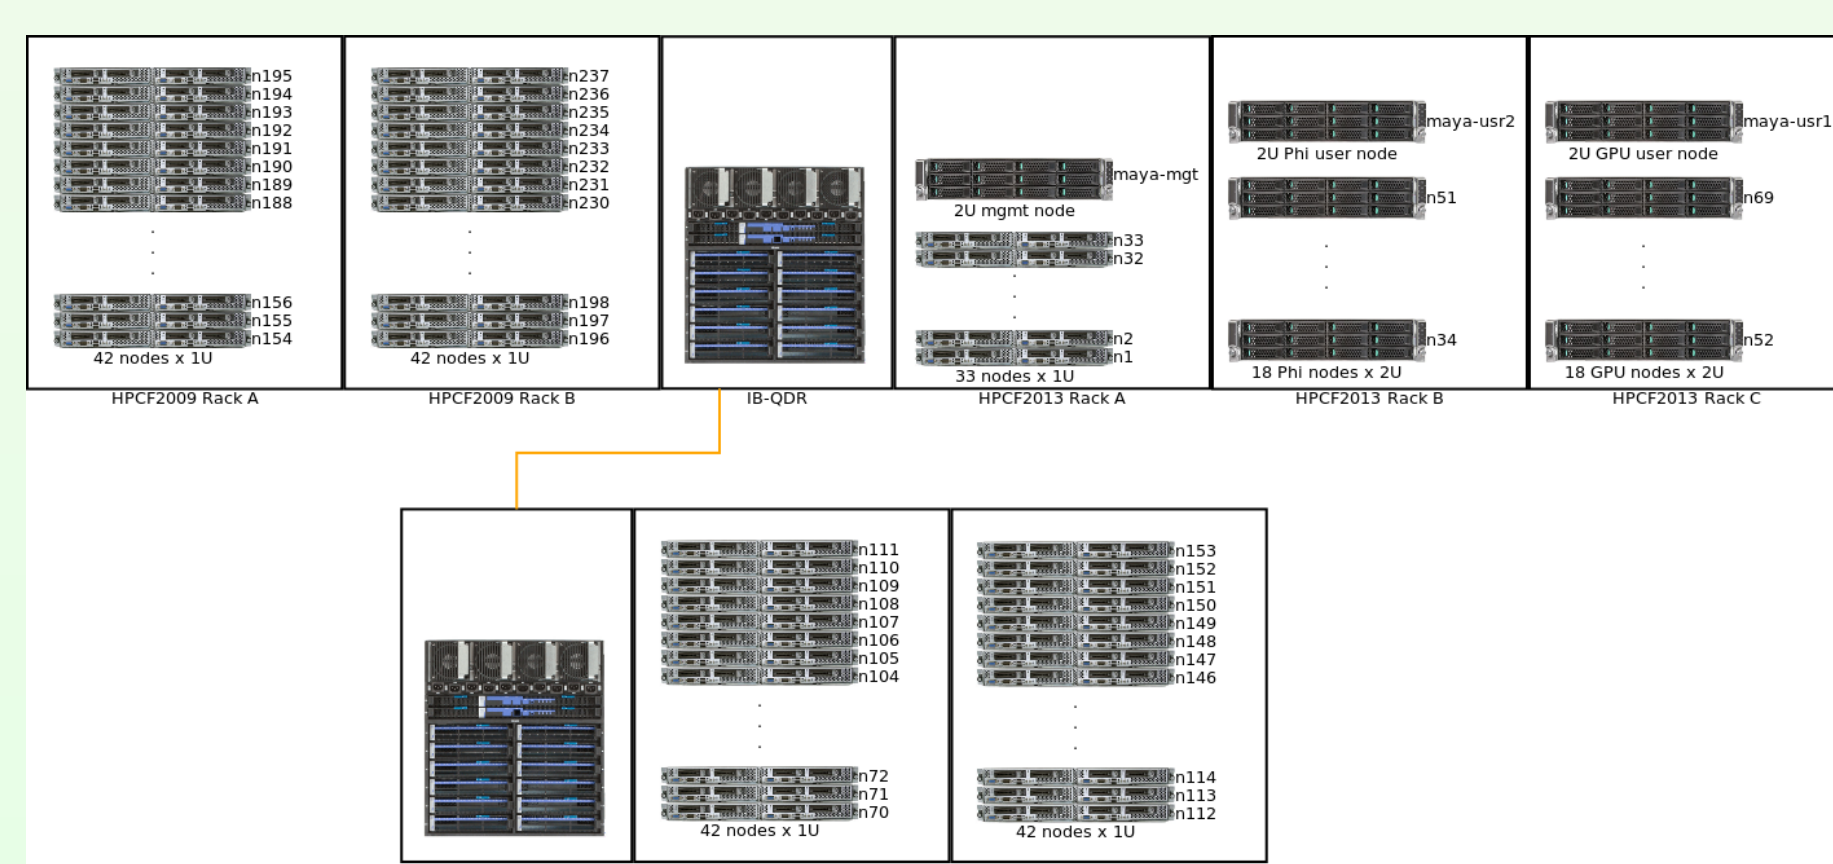
## 240-Node Cluster maya

- **HPCF2013 = maya (2013) = \$540,000:**  
72 nodes, each with two 2.6 GHz eight-core Intel E5-2650v2 Ivy Bridge CPUs:
  - 34 CPU-only nodes
  - 19 hybrid nodes with two NVIDIA K20 GPU
  - 19 hybrid nodes with two Intel Phi 5110P
- **HPCF2010 = maya (2010) = gift from NASA:**  
84 nodes, each with two 2.8 GHz quad-core Intel Nehalem X5560 CPUs
- **HPCF2009 = maya (2009) = \$600,000:**  
84 nodes, each with two 2.6 GHz quad-core Intel Nehalem X5550 CPUs

Networks connecting all components:

- quad-data rate (QDR) InfiniBand interconnect for HPCF2013 and HPCF2009
- dual-data rate (DDR) InfiniBand interconnect for HPCF2010

Storage of more than 750 TB connected by IB.

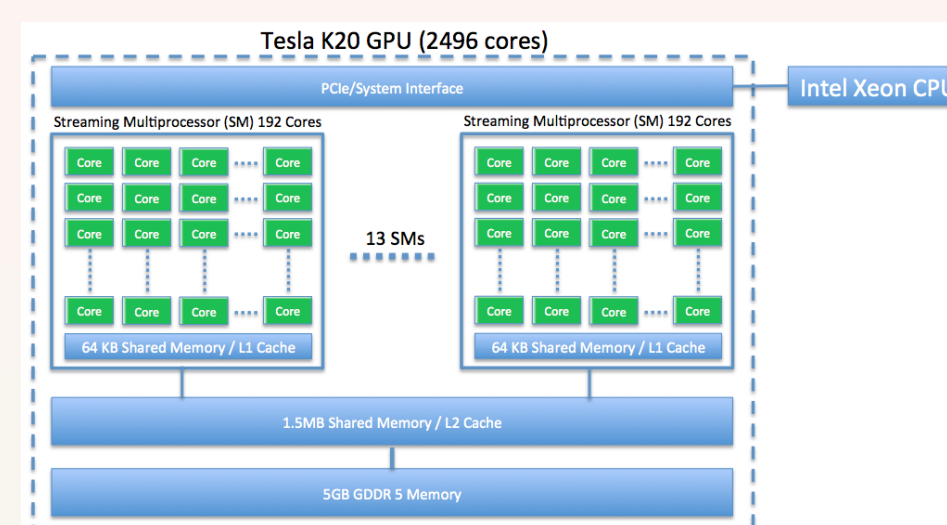


## Calcium Induced Calcium Release (CICR) Solved with CPUs Only [4]

Wall clock time in HH:MM:SS of the CICR problem solved with first order finite volume method [1] on HPCF2013 by number of nodes and MPI processes per node. Mesh resolution  $N_x \times N_y \times N_z = 128 \times 128 \times 512$ , DOF = 25,610,499. ET indicates "excessive time required". N/A indicates that the case is not feasible due to  $p > (N_z + 1)$ .

	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes	64 nodes
1 process per node	ET	ET	69:15:37	34:51:02	17:31:44	08:59:06	04:49:17
2 processes per node	ET	69:46:29	35:16:03	17:45:06	09:00:47	04:47:14	02:43:04
4 processes per node	72:31:51	36:34:34	18:36:29	09:32:04	05:01:44	02:50:34	01:46:47
8 processes per node	42:01:27	26:23:03	11:03:41	05:46:44	03:09:23	01:56:47	01:23:57
16 processes per node	26:53:37	13:56:38	07:21:17	03:54:47	02:17:48	01:40:35	N/A

## Hybrid Nodes with 2 GPUs [2]

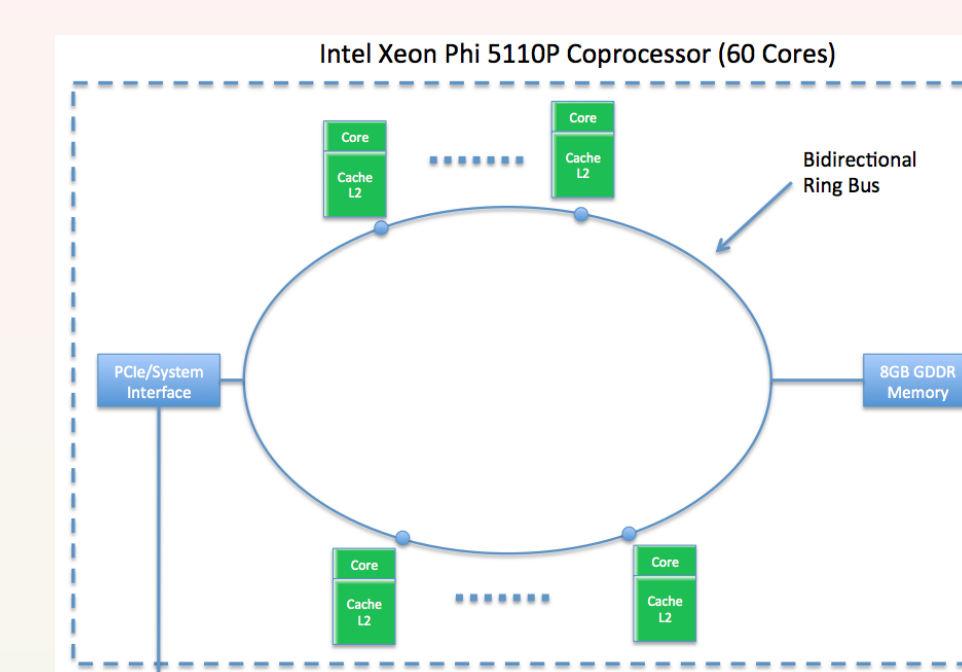


19 hybrid nodes contain two NVIDIA K20 GPUs with 2496 computational cores with 5 GB of global memory.

Wall clock time (speedup) of CICR problem against one 16-core node with CUDA+MPI code.

nodes (GPU/node)	$128 \times 128 \times 512$
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)

## Hybrid Nodes with 2 Intel Phi



19 hybrid nodes contain two 60-core Intel Phi 5110P. The 8 GB of memory are connected through a bidirectional ring bus.

Intel Phi as accelerators can be used to increase throughput, and the x86 compatible architecture can minimize the programming effort. There are three ways to access Intel Phi: native execution on the Phi, offloading to the Phi, and symmetric mode using both CPU and Phi heterogeneously. Wall clock time in HH:MM:SS of CICR problem on one Phi in native mode with original MPI code.

MPI processes	$64 \times 64 \times 256$	$128 \times 128 \times 512$
1 process	ET	ET
2 processes	ET	ET
4 processes	88:50:37	ET
8 processes	43:39:36	ET
16 processes	24:15:25	ET
32 processes	12:09:55	ET
64 processes	07:36:25	ET
128 processes	07:01:37	66:29:19

## Acknowledgments and References

- UMBC, CIRC, REU
- NSF (MRI, SCREMS), NSA, NASA

- [1] Schäfer, Huang, et al., *NMPDE*, 2015  
 [2] Huang, Ph.D. Applied Mathematics, 2015  
 [3] REU Site: HPCF-2013-13, HPCF-2014-14  
 [4] HPCF-2015-6, HPCF-2015-7, HPCF-2015-8