Performance Comparison of Intel Xeon Phi Knights Landing

UMBC REU Site: Interdisciplinary Program in High Performance Computing
Ishmail A. Jabbie[1], George Owen[2], Benjamin Whiteley[3],
RA: Jonathan S. Graf[1], Faculty mentor: Matthias K. Gobbert[1], Client: Samuel Khuvis[4]
[1]UMBC, [2]Louisiana State University, [3]University of Maryland, Eastern Shore, [4]ParaTools, Inc.

Motivation

The Intel Xeon Phi is a many-core processor family with theoretical peak performance of over 1 TFLOP/s in double precision, significantly better than even modern multi-core CPUs. A test problem in C using MPI and OpenMP compares the performance of the first and second generations of the Phi, codenamed Knights Corner (KNC) from 2013 and Knights Landing (KNL) from 2016, respectively, as well as contrasts the performance of the two different memories available on the KNL.

Test Problem

Classical elliptic test problem: Poisson equation with homogeneous Dirichlet boundary conditions

$$\begin{aligned} -\Delta u &= f \qquad \text{in } \Omega, \\ u &= 0 \qquad \text{on } \partial\Omega, \end{aligned}$$

on domain $\Omega = (0,1) \times (0,1) \subset \mathbb{R}^2$.
The equation is discretized by the finite difference method and the resulting system of linear equations solved by the conjugate gradient method.
The numerical method is parallelized in C with MPI and OpenMP.

References and Acknowledgments

[1] For more information on the Intel Xeon Phi, `https://software.intel.com/en-us/xeon-phi/x200-processor`

[2] Samuel Khuvis, Ph.D. Thesis, Applied Mathematics, UMBC, 2016

[3] Full technical report: HPCF–2016–16 `hpcf.umbc.edu` > Publications.

- REU Site: `hpcreu.umbc.edu`
- NSF, NSA, DOD, UMBC, HPCF, CIRC
- TACC, XSEDE: `www.xsede.org`
- U. of Oregon: `nic.uoregon.edu/prl`

Knights Corner (KNC)

First-generation Phi from 2013 — KNC SE10P in Stampede at TACC:

- 61 cores *[contrast: CPU has 8 or 16 cores]*
- 1 VPU up to 8 double additions per core
- 8 GB of on-board GDDR5:
  form of RAM adapted to perform better with graphical data
- Bi-directional ring bus
- Linux Micro-OS
- Co-processor only

Intel compiler suite used on all systems.

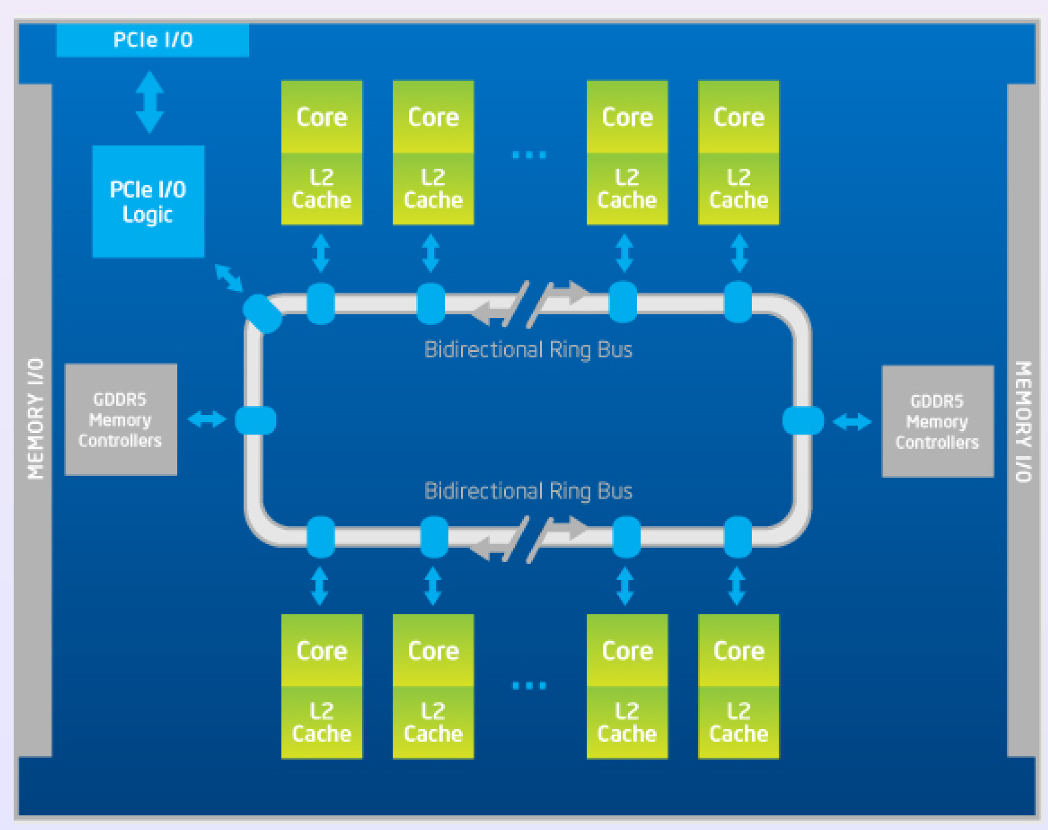

Knights Landing (KNL)

Second-generation Phi from 2016 — KNL 7250 in PRL at U. of Oregon:

- 68 cores
- 2 VPUs up to 16 double additions per core
- **16 GB on-board MCDRAM:**
  High performance 3D RAM, much faster than GDDR5 or DDR4, designed for high bandwidth memory
- Server contains 98 GB of DDR4 RAM:
  Larger but slower system memory
- **2D mesh network**
- Full Linux-based OS
- Full stand alone processor

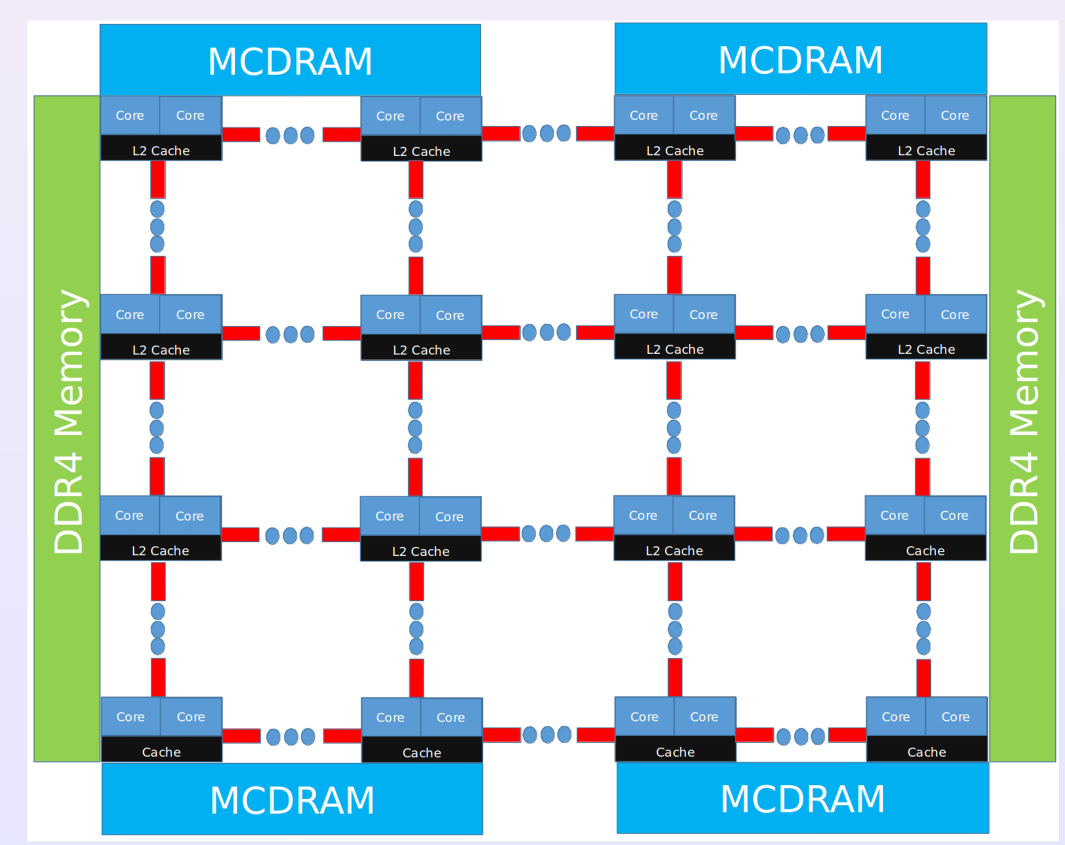

Solution of Test Problem on $N \times N = 8{,}192 \times 8{,}192$ Mesh

KNC at TACC with GDDR5 RAM using 240 threads [1]: observed wall clock time in MM:SS

MPI proc	1	2	4	8	15	16	30	60	120	240
Threads	240	120	60	30	16	15	8	4	2	1
GDDR5	28:24	28:20	27:51	23:08	23:06	23:00	22:24	22:45	22:43	25:37

KNL at Oregon with DDR4 and MCDRAM using 272 threads: observed wall clock time in MM:SS

MPI proc	1	2	4	8	16	17	34	68	136	272
Threads	272	136	68	34	17	16	8	4	2	1
DDR4	26:02	25:07	24:38	24:25	24:24	36:29	37:40	37:54	39:06	41:00
MCDRAM	05:49	05:43	05:39	05:35	05:36	08:22	08:49	08:41	08:37	08:57

Conclusions

- **The KNL using MCDRAM is dramatically faster in all cases.**
- Despite DDR4 being a slower form of memory, KNL using DDR4 is comparable in most cases to KNC using GDDR5.
- **For both MCDRAM and DDR4 on the KNL, using more threads than MPI processes is significantly faster than the inverse.**
- KNL distributes cores optimally to use resources and channels of the system. Threads allow the processor to assign the cores in order, while MPI assigns processes randomly.