

Lawrence Berkeley National Laboratory

LBL Publications

Title

Porting the MIMD Lattice Computation (MILC) Code to the Intel Xeon Phi Knights Landing Processor

Permalink

<https://escholarship.org/uc/item/8026s7t8>

Authors

Li, R
DeTar, C
Doerfler, D
et al.

Publication Date

2016-06-23

Peer reviewed

Porting the MIMD Lattice Computation (MILC) Code to the Intel Xeon Phi Knights Landing Processor

Ruizi Li¹, Dhiraj Kalamkar², Ashish Jha², Steven Gottlieb¹, Carleton
DeTar³, Doug Toussaint⁴, Balint Joo⁵, and **Douglas Doerfler**⁶

¹Indiana University, ²Intel Corp.,
³University Of Utah, ⁴University of Arizona,
⁵Thomas Jefferson National Accelerator Facility,
⁶Lawrence Berkeley National Laboratory

Introduction

- **The MILC code was chosen by NERSC to be a tier 1 application in its NESAP (NERSC Exascale Science Applications Program) collaborative**, an effort in which NERSC will partner with code teams to prepare for the Cori supercomputer
- NESAP provides access to staff from NERSC, Cray and Intel and also provides unique resources such as early hardware test beds and advanced training
- In addition, **an Intel Parallel Computing Center has been established at Indiana University** (in part) to help port the MILC code to the Intel Xeon Phi processor


Target Hardware Architecture

- Intel Endeavor Cluster
 - Multiple node types
 - OmniPath high speed interconnect
- Intel KNL
 - KNL preproduction, B0 stepping
 - 64 cores @ 1.3 GHz, 4 hyper-threads/core
 - 16 GB MCDRAM (>460 GB/s peak BW)
 - 96 (6x16) GB DDR4 @ 2133 MHz (102 GB/s peak BW)
- Intel Broadwell
 - Dual socket, 18 cores/socket @ 2.30 GHz
 - 128 GB (8x16) DDR4 @ 2400 MHz (153.6 GB/s peak BW)

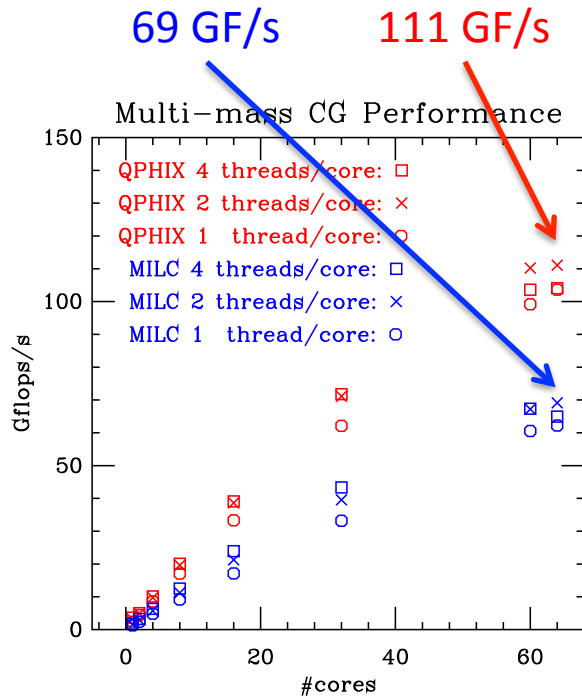
The QPhiX Library

- An open source library for generating vectorized lattice QCD code on multi-core architectures
- Originally developed for *Wilson* and *Wilson Clover* quarks on Knights Corner
- **MILC effort adapts QPhiX to *Staggered Fermions***
- Library is based on code generation and intrinsics
 - Has been targeted to SSE, AVX, AVX2, AVX512 and BlueGene/Q QPX
- **Major feature is improved vectorization by blocking in the X-Y planes of the lattice**
- NESAP effort has also provided further OpenMP optimizations
- **The DSLASH operator is implemented with QPhiX, *but our analysis is for the entire Multi-mass CG solve***


Single Node Weak Scaling



MCDRAM Flat




MCDRAM Cache



DDR only

Multi-node Weak Scaling



Thank You

dwdoerf@lbl.gov