Twisted-Mass Lattice QCD using OpenCL

An Update

Matthias Bach

Frankfurt Institute for Advanced Studies

Institut für Informatik

in collaboration with:
C. Pinke, O. Philipsen, V. Lindenstruth

31st International Symposium on
Lattice Field Theory
Mainz, 30 July 2013

LOEWE-CSC

- Frankfurt University
- 786 GPU Nodes
- Node:
  - 2 AMD Magny-Cours
  - AMD Radeon HD 5870
- QDR Infiniband
- 285 TFLOPS (#22)
- 740.78 MFLOPS/W
- Green500 #8 Nov 10

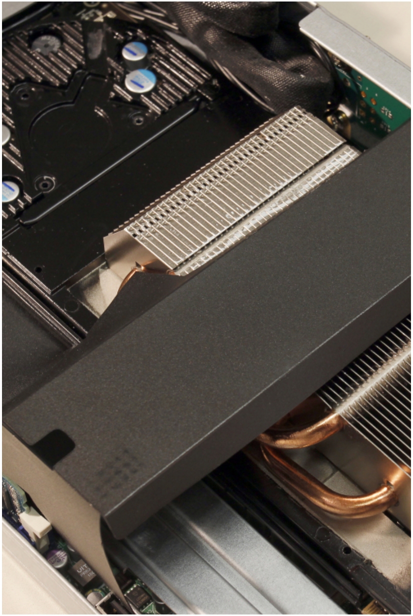

Sanam

- Coop:
  - FIAS (Frankfurt)
  - KACST (Saudi-Arabia)
- 300 Nodes
- Node:
  - 2 Intel Xeon E5-2650
  - 2 AMD FirePro S10000
- FDR Infiniband
- 532 TFLOPS (#52)
- 2351 MFLOPS/W
- Green500 #2 Nov 12

OpenCL and CL^{2}QCD

OpenCL

- Open Standard
- Like CUDA Driver API
  - Kernels
  - Threads
  - C-API
- Wide hardware support
  - CPUs
  - AMD GPUs
  - NVIDIA GPUs
  - Xeon Phi
  - Mobile Devices

CL^{2}QCD

- Twisted Mass / Pure Wilson
- All calculations in OpenCL
- Layered
  - Hardware
  - Kernels
  - QCD-Types
  - Algorithms
- Seperation of Concerns
  - Hardware Specifics
  - Optimization
  - Application Logic

Double-Precision $\not{D}$ Performance

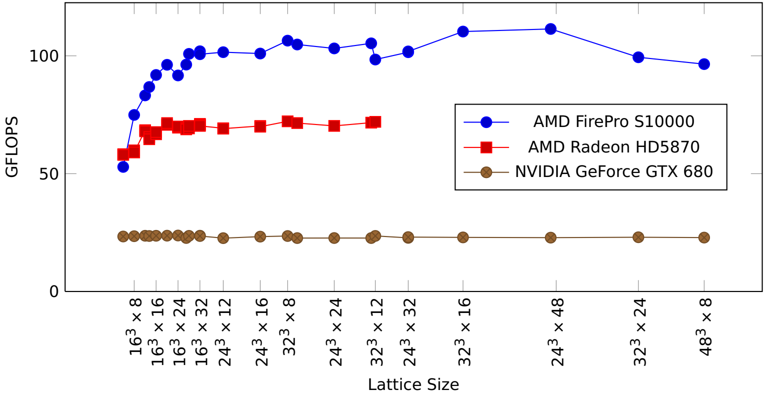

See also: *M. Bach, V. Lindenstruth, O. Philipsen, and C. Pinke, "Lattice QCD based on OpenCL," Computer Physics Communications, p. 19, Mar. 2013.*

CG Performance – Weak Scaling

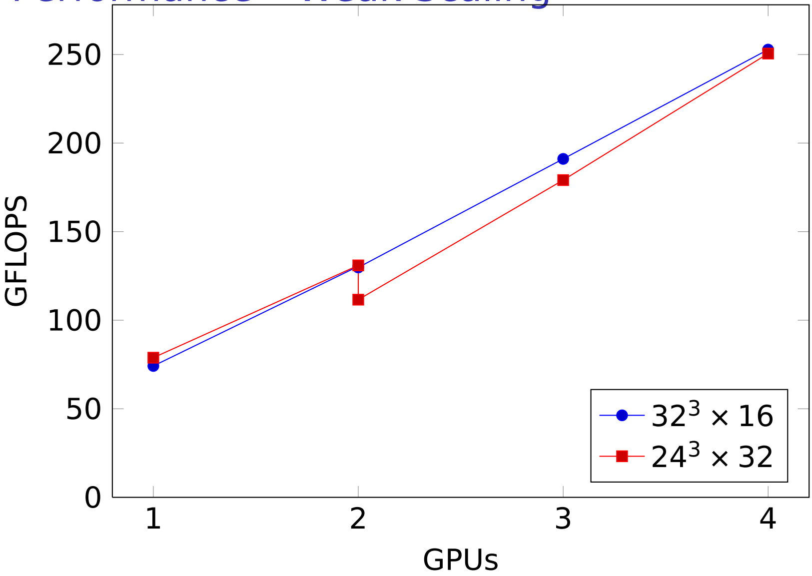

Constant lattice size per GPU

CG Performance – Hard Scaling

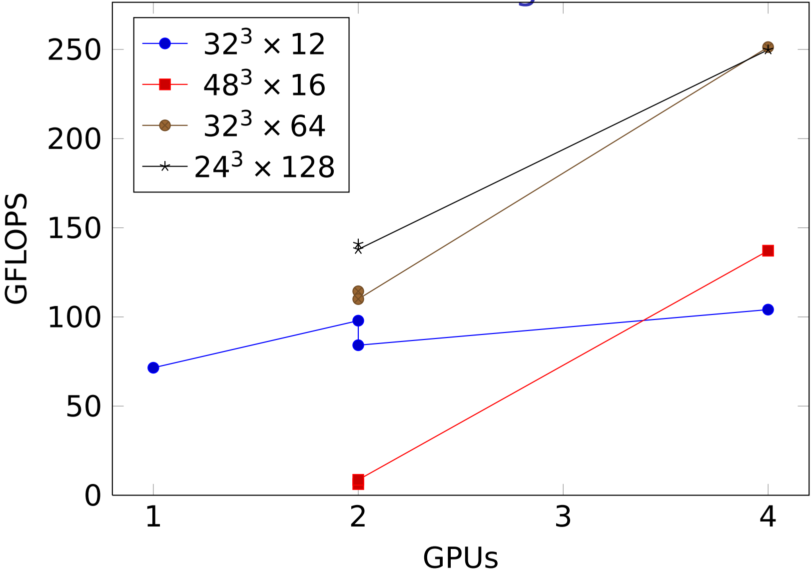

Varying lattice size per GPU

Average Power Consumption

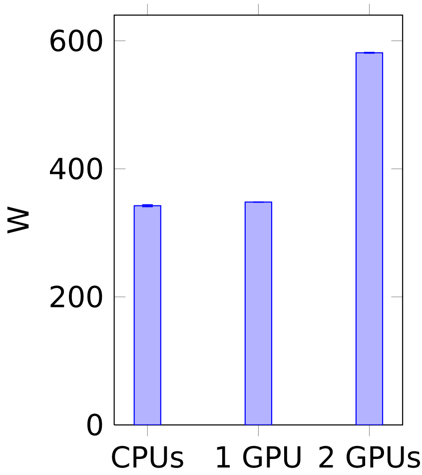

CPU System:

- 2 AMD Bulldozer
  8 Cores, 3 GHz
- *tmlqcd*

  K. Jansen and C. Urbach, “tmLQCD: a program suite to simulate Wilson Twisted mass Lattice QCD,” Quantum, May 2009, pp. 1–44, 2009.

GPU System:

- 2 AMD Magny-Cours
  12 Cores, 2 GHz
- 1 to 2 AMD Radeon
  HD7970
- *CL²QCD*

Energy-Efficiency

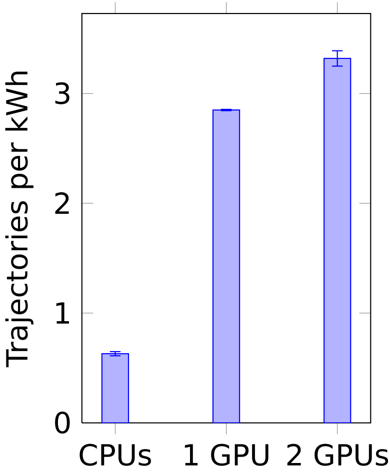

HMC Setup:

- $32^3 \times 12$
- $N_f = 2$ Twisted-Mass Wilson Fermions
- $m_\pi \approx 270\,\text{MeV}$

Multiple GPUs:

- 2nd GPU: Seperate HMC Chain
- Projection to 8 GPUs: 3.9 to 4.25 trajectories per kWh

Total Cost of Aquisition in General Purpose Systems

HMC in Sanam:

- GPU Throughput:
  16 times CPU
- GPU Cost:
  $\approx 25\,\%$
- GPU Cost Advantage:
  ≈ 12

Total Cost of Aquisition in General Purpose Systems

HMC in Sanam:

- GPU Throughput:
  16 times CPU
- GPU Cost:
  ≈ 25 %
- GPU Cost Advantage:
  ≈ 12

Cost for Inversions:

- Kraken (x86)
  1.14 $/MFLOPS
- Sequoia (BlueGene/Q)
  0.040 $/MFLOPS
- Sanam (x86 + GPU)
  0.038 $/MFLOPS

Cost on Kraken and Sequoia estimated based on publications.

Total Cost of Aquisition in Dedicated Systems

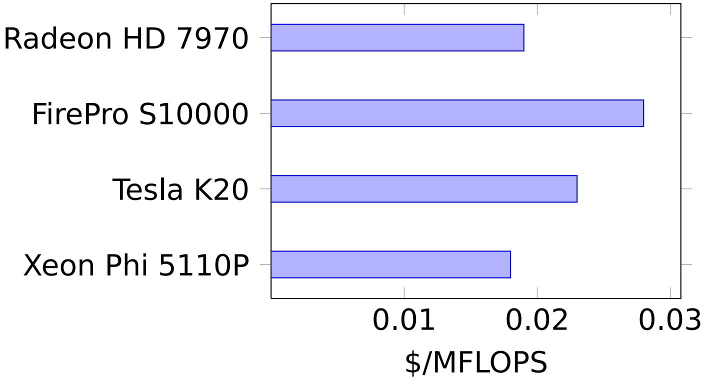

Based on retail prices

Total Cost of Aquisition in Dedicated Systems

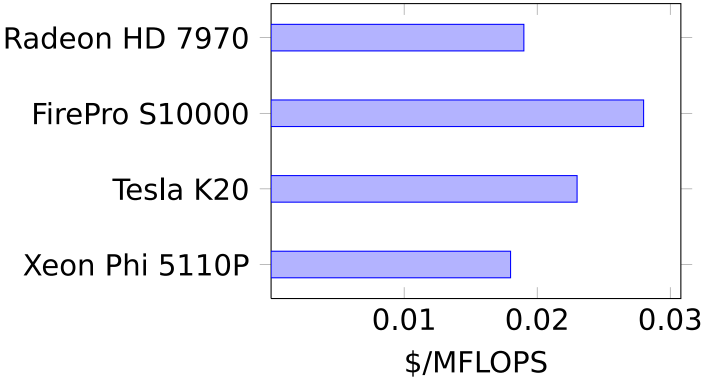

Based on retail prices

Comparision should use HMC trajectories!

Application

C. Pinke and O. Philipsen,
"The nature of the Roberge-Weiss transition in $N_f = 2$ QCD with Wilson fermions",
Monday, 17:50

- Utilized LOEWE-CSC and Sanam

Outlook

Achieved:

- 100 GFLOPS DP $D\!\!\!/$ on AMD GPU
- Works on NVIDIA / Intel
- Full HMC on GPU
- Scaling on GPUs within Node
- Four Times as Energy-Efficent as CPU
- Cost-Efficient

In Progress / TODO:

- Staggered Fermions
- Mixed-Precision Solvers
- MPI
- NVIDIA Optimizations
- Intel Optimizations