

LQCD software (DWF/Mobius) optimization (BG/Q & GPU)



COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK

Chulwoo Jung, Hyung-Jin Kim
Brookhaven National Laboratory

April 17, 2014
USQCD SciDAC 2014 Software Workshop
Jefferson Lab, Newport News, Virginia

G-parity expansion of CPS/BFM: Mixed Precision Multimass Solver (C.Kelly)


$$\hat{G} = \hat{C} e^{i\pi \hat{I}_y}, \hat{G} \begin{pmatrix} u \\ d \end{pmatrix} = \begin{pmatrix} -C \bar{d}^T \\ C \bar{u}^T \end{pmatrix} \quad C = \gamma^2 \gamma^4$$
$$\hat{G} |\pi^\pm\rangle = -|\pi^\pm\rangle \quad \hat{G} |\pi^0\rangle = -|\pi^0\rangle$$

Kim & Christ arXiv:hep-lat/0210003,


- Need new ensembles with G-parity.
- Inherently 2-flavor: doubles flavors.
- Strange quark simulated as s/s' .
- Normal Wilson/DWF/Mobius would give 4 flavors instead of 2 \rightarrow light flavors needs to be simulated with RHMC

- Light quark makes the range of eigenvalues for rational approx. larger (lowest ev $0.1 \rightarrow 0.0002$)
- Significant difference in performance between SP and DP (23.9 vs. 9.7 Gflops for CG $V = 8^4 \times 12(\times 2)$)


SP vs. DP precision true residuals of the lightest pole($\sim 10^{-5}$)




Replace SP residual with DP once per 100 iter.



x and p in DP, r in SP.



Reset r in DP once per 4000 iter



Current strategy:

- Multimass with x, p in DP, r in SP
- Finish each pole with shifted defect correction.
- Improved total RHMC time by $\sim 40\%$.
- Precision switching, etc can be further optimized.
- Defect correction for 2nd,.. poles can still be significant.
- For DWF/Mobius, 1-flavor without RHMC may be possible. arXiv:1403.1683
- Further optimization of Implicitly restarted Lanczos using XLC intrinsics(D. Zhang, Similar to Alexei's work on EigCG). Now Dslash is dominating Lanczos.
- HDCG(arXiv:1402.2585): with HDCG?
- BNL institutional cluster (GPU & Phi)

Lanczos Performance 1 (Preliminary)

- **Scalability Test** | 16x16x16x32x4 DWF lattice, k=5, p=5, unit: second, tested on K20x@BNL

Nodes→	1	2	4
Including Total IO	12.55 (1,1,1,1)	6.70 (1,1,1,2)	3.60 (1,1,1,4)
Dirac kernel	0.036(?)	3.79	2.26
Scaling factor	1	1.87	3.48

- **GPU architecture Test** | same condition, tested on C2050, K20x, K40s@BNL

	C2050 E5620 2.4Ghz	K20x E5-2670 2.6Ghz	K40s I5-3570 3.4Ghz
2 node total(sec)	380	287	158
Scaling factor	1	1.32	2.40

Lanczos Performance 2(Preliminary)

- CPU vs GPU | 16x16x16x32x16 DWF lattice, k=20, p =10, tested on k40s@BNL

	2 x k40s 15 3570 3.4 Ghz	4 x 8 core AMD 6128 2.0Ghz
Total time(sec)	1433 (1,1,1,2)	3011 (2,2,2,4)
Scaling factor	2.10	1

CG Performance

- CG performance on QUDA

@ K20m	4	8	16	32
24^3x64x16	892 (1,1,1,4)	1110 (1,1,1,8)	1310 (1,2,2,4)	1692 (1,1,4,8)
32^3x64x16	N/A	1110 (1,1,2,4)	1442 (1,2,2,4)	1743 (2,2,2,4)

- CG performance on various platforms

24x24x24x64	RICC (Fermi C2075)	FNAL (Fermi M2050)	Jlab (Kepler K20m)	BNL K40s
2 node(s=8)	232		502.5	599
4 node(s=16)	399		892	
8 node(s=16)		670		

- CG performance in 48x48x48x96 lattice (~300GB mem used)

Node geometry	72(2,3,3,4)	72(1,3,3,8)
48x48x48x96	3547	3617

Performance(DWF and Mobius DWF)

CG performance (DWF vs MDWF) (GFLOPS for all GPUs, Single-Double mixed precision)

24x24x24x64x16 (K40, BNL)	MDWF	DWF
2 node	553	596

MDWF CG performance (Fermi Vs Kepler) (GFLOPS for all GPUs, Single-Double mixed precision)

24x24x24x64 x8	C2050	K20m	K20x
4 node	390	840	910

Performance(DWF and Mobius DWF)

MADWF Iteration time sequence

24x24x24x64x16 (K20x, BNL)	Time (sec)	Total time (sec)
Sninor reordering(CPU)	3.31 x 30	~99.3
EO Preconditioning (CPU)	43.1 x 15	~646
MDWF CG(GPU,Iter = 7941)	622	622
Restore Full vector(CPU)	19.0 x 15	~285
Spin Projection(CPU)	3.85 x 12	46.2
		= ~ 1700
Actual running time		1461