

TWQCD's dynamical DWF project

**TWQCD Collaboration: Ting-Wai Chiu^{*,1,2†}, Tian-Shin Guu³, Tung-Han Hsieh⁴,
Chao-Hsi Huang³, Yu-Ying Lee¹, Yao-Yuan Mao¹, Kenji Ogawa¹, Po-Kai Tseng²**

¹ *Department of Physics, and Center for Theoretical Sciences, National Taiwan University, Taipei 10617, Taiwan*

² *Center for Quantum Science and Engineering, National Taiwan University, Taipei 10617, Taiwan*

³ *Center for General Education, and Institute of CSIE, National ILan University, I-Lan 260, Taiwan*

⁴ *Research Center for Applied Sciences, Academia Sinica, Taipei 115, Taiwan*

We present an overview of our project of simulation of unquenched lattice QCD with optimal domain-wall quarks, using a GPU cluster currently constituting of 16 units of Nvidia Tesla S1070 plus 64 graphic cards with Nvidia GTX285 (total 128 GPUs with 128 Teraflops peak), attaining sustained computing power of 15.36 Teraflops. The first production run in two-flavor QCD is ongoing, using the Iwasaki gauge action on a set of lattices with sizes $16^3 \times (32, 10, 8, 6, 4) \times (16, 32)$ at the lattice spacing $a \sim 0.1$ fm, with eight sea quark masses down to $m_\pi \simeq 200$ MeV. We outline our simulation algorithm, and describe the present status of the production run. Preliminary results of pseudoscalar mass and decay constant are also presented.

The XXVII International Symposium on Lattice Field Theory - LAT2009

July 26-31 2009

Peking University, Beijing, China

*Speaker.

†Email: twchiu@phys.ntu.edu.tw

1. Introduction

Simulation of unquenched lattice QCD with exact chiral symmetry is a grand challenge among all sciences. Even for a modest $16^3 \times 32$ lattice with lattice spacing $a \sim 0.1$ fm, it often requires a supercomputer with peak computing power more than 50 Teraflops (e.g., 10 racks of IBM BlueGene/L). Therefore, only 2-3 lattice QCD groups around the world could afford to perform the simulation of unquenched lattice QCD with the domain-wall fermion [1], or the overlap-Dirac fermion [2]. (For a recent review of generating gauge ensembles of QCD, see Ref. [3].)

However, this scenario has been undergoing a dramatic change during the last 9 months. With the emergence of low-cost and computationally powerful Graphic Processing Unit (GPU), now plugging a graphic card with Nvidia GTX285 (240 cores, one Teraflops peak) into a PC immediately turns the system into a powerful device, attaining a sustained 140 Gflops for the conjugate gradient (CG) with mixed precision. This implies that one can perform the Hybrid Monte Carlo (HMC) simulation of unquenched lattice QCD with exact chiral symmetry, using a single PC with a GTX285 graphic card, rather than a cluster of 128 nodes. Moreover, the cost of such a system (PC + GPU card) could be less than \$600. In terms of price/performance, a GPU cluster is only 1/100 of a supercomputer (e.g., IBM BlueGene/L). Also, the power consumption of a GPU cluster is less than 1/10 of any supercomputer, for the same amount of (sustained) Teraflops. Evidently, these salient features of GPU have revolutionary impacts to lattice QCD as well as all computational sciences. For a recent review of QCD on GPUs, see Ref. [4].

The dynamical DWF project of TWQCD Collaboration was started around November 2007, soon after CUDA version 1.1 became available¹. This was the right time for us to cut in since we could write our QCD code in CUDA rather than the much more tedious Cg [5]. We completed our first CUDA code around March 2008, which implemented the nested conjugate gradient with mixed precision to compute the overlap fermion propagator, in which the inner CG (i.e., the sign function times a vector) was computed with the two pass algorithm [6, 7]. Note that Nvidia GTX280 had not been available at that time, thus all double precision operations were performed by the CPU, yet we still achieved 20 Gflops (sustained) using the Nvidia GTX8800 + Intel Q6600 (2.40 GHz).

At lattice 2008, two groups reported using GPU for the Wilson-Dirac matrix multiplication with a vector [8, 9], attaining ~ 100 Gflops with Nvidia GTX280. On the other hand, our aim had been to simulate unquenched lattice QCD with exact chiral symmetry, in which the major difficulty is how to preserve chiral symmetry to a high precision and also to sample all topological sectors ergodically. For lattice QCD with the conventional domain-wall fermion [11], it suffers from a large chiral symmetry breaking (i.e., large residual mass) for $N_s = 16$, which becomes problematic for finite temperature QCD. On the other hand, for lattice QCD with overlap-Dirac fermion, the action is discontinuous at the boundary between different topological sectors, and it is very costly to tunnel through the topological boundary using the refraction/reflection algorithm [10]. A plausible way to proceed is to fix the topology [12], then use the measured values of topological susceptibility to remove the artifacts in any physical observables due to the fixed topology in a finite volume. However, it has not been proved that suppressing the near zero modes (with the unphysical Wilson fermions and the associated twisted-mass ghosts) does not violate the ergodicity in a fixed topological sector.

¹T.W.C. thanks Kelvin Chiu (then at Intel Corporation, CA) for his helpful suggestion.

Now we turn to the optimal domain-wall fermion. Unlike the conventional domain-wall fermion suffering from large chiral symmetry breaking, the optimal domain-wall fermion attains the (mathematically) maximal chiral symmetry by judiciously fixing the weight ω_s for each layer in the fifth dimension, according to the formula derived in Ref. [13]. Then the effective 4D Dirac operator is exactly equal to the overlap-Dirac operator with the sign function approximated by the Zolotarev optimal rational polynomial.

The first step of our project was to develop the CUDA code for the CG of the 5D optimal domain-wall quark matrix, the most computationally intensive part in HMC. After tuning our CUDA code for 5 months, in January 2009, we achieved 120 Gflops for a single GTX280.

The outline of this paper is as follows. In Section 2, we describe the hardware of our project. In Section 3, we outline our HMC simulation with the optimal domain-wall fermion. In Section 4, we briefly describe the GPU optimization of our CG code. In Section 5, we describe the present status of our production runs. In Section 6, we present preliminary results of pseudoscalar mass and decay constant. Finally, in Section 7, we conclude with some remarks.

2. The hardware

Our hardware is a GPU cluster currently constituting of 16 units of Nvidia Tesla S1070 plus 64 Nvidia GTX285 graphic cards. Each unit of Nvidia Tesla S1070 contains 4 Tesla GPUs, each having 4 Gbyte RAM, while each GTX285 has 1 or 2 Gbyte RAM. Each unit of Tesla S1070 is connected to the PCI Express x16 slots of a Tyan server through two interconnect cables and two interface cards. Each Tyan Server has two Intel Quad-Core Xeon processors and 32 GB RAM. For the GTX285 cards, every two GTX285 cards are plugged into the PCI Express x16 slots of a 2U server with one Intel i7 CPU and 12 Gbyte RAM. The hard disk storage of our system is over 300 Terabytes, which is managed through the Lustre cluster file system (<http://www.sun.com/software/products/lustre/index.xml>). The hardware cost of our system is around \$220,000. The sustained computing power of our system attains 15.36 Teraflops for HMC simulation of lattice QCD with optimal domain-wall quarks. We use the Condor to manage the workload of our computing jobs (see <http://www.cs.wisc.edu/condor/>).

3. Hybrid Monte Carlo simulation with optimal domain-wall quark

The action of optimal domain-wall fermion [13] is defined as

$$\begin{aligned} S_{odwf} &= \sum_{s,s'=1}^{N_s} \sum_{x,x'} \bar{\psi}_{xs} \{ (\omega_s D_w + 1)_{xx'} \delta_{ss'} + (\omega_s D_w - 1)_{xx'} (P_+ \delta_{s',s-1} + P_- \delta_{s',s+1}) \} \psi_{x's'} \quad (3.1) \\ &= \bar{\Psi} D_{odwf} \Psi \end{aligned}$$

with boundary conditions $P_+ \psi_{x,0} = -(m_q/2m_0) P_+ \psi_{x,N_s}$, $P_- \psi_{x,N_s+1} = -(m_q/2m_0) P_- \psi_{x,1}$, where $P_{\pm} = (1 \pm \gamma_5)/2$. Here D_w is the standard Wilson Dirac operator plus a negative parameter $-m_0$ ($0 < m_0 < 2$), m_q is the bare quark mass, and the weights $\{\omega_s\}$ are specified by the exact formula derived in Ref. [13]. The optimal domain-wall fermion operator D_{odwf} can be transformed to

$$D_{opt} = D_w + M, \quad M = \frac{1}{\sqrt{\omega}} \frac{1+L}{1-L} \frac{1}{\sqrt{\omega}}, \quad L = P_+ L_+ + P_- L_-, \quad (3.2)$$

where the ω denotes the diagonal matrix of the weights $\{\omega_s, s = 1, \dots, N_s\}$, and

$$(L_+)_{s,s'} = \begin{cases} -\delta_{s',s-1}, & 1 < s \leq N_s, \\ (m_q/2m_0)\delta_{s',N_s}, & s = 1, \end{cases}, \quad L_- = (L_+)^\dagger. \quad (3.3)$$

Obviously, M does not carry the color and 4D space-time indices, and it can be solved analytically. Since $\det(D_{opt})$ only differs from $\det(D_{odwf})$ by a factor independent of the gauge field, we can use D_{opt} to perform the HMC simulation rather than D_{odwf} .

With even-odd preconditioning in the 4D space-time lattice, D_{opt} can be written as

$$D_{opt} = \begin{pmatrix} X & D_w^{eo} \\ D_w^{oe} & X \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ D_w^{oe} X^{-1} & 1 \end{pmatrix} \begin{pmatrix} X & 0 \\ 0 & X - D_w^{oe} X^{-1} D_w^{eo} \end{pmatrix} \begin{pmatrix} 1 & X^{-1} D_w^{eo} \\ 0 & 1 \end{pmatrix} \quad (3.4)$$

where $X = 4 - m_0 + M$. Since $\det D_{opt} = \det(1 - D_w^{oe} X^{-1} D_w^{eo} X^{-1}) \det X$, and X does not depend on the gauge field, we can use the operator

$$C = 1 - D_w^{oe} X^{-1} D_w^{eo} X^{-1} \quad (3.5)$$

for the HMC simulation. After including the Pauli-Villars fields (with $m_q = 2m_0$), the pseudo-fermion action for 2-flavor QCD ($m_u = m_d$) can be written as

$$S_{pf} = \phi^\dagger C_{PV}^\dagger (CC^\dagger)^{-1} C_{PV} \phi, \quad (3.6)$$

where $C_{PV} = C(m_q = 2m_0)$.

In the HMC simulation, we first generate random noise vector ξ with Gaussian distribution, then we obtain $\phi = C_{PV}^{-1} C \xi \equiv D \xi$ with CG. With fixed ϕ , the system is evolved with a fictitious Hamiltonian dynamics, the so-called molecular dynamics (MD). In the MD, we use the Omelyan integrator [14], and the Sexton-Weingarten multiple-time scale method [15]. The most time-consuming part in the MD is to compute $\eta = (DD^\dagger)^{-1} \phi$ with CG, which is required for the evaluation of the fermion force in the equation of motion of the conjugate momentum of the gauge field. Thus we program the GPU (Nvidia T10 or GTX285) to compute $\eta = (DD^\dagger)^{-1} \phi$, using CG with mixed precision. Also, we have ported the computation of the gauge force and the gauge field update to the GPU. Our next step is to implement the computation of the fermion force entirely inside GPU. In some of our HMC simulations, we also introduce an extra heavy fermion field to reduce the condition number of the CG involving the light quark [16]. In this case, we can have both CPU and GPU compute concurrently, e.g., while the GPU is working on the CG of the light quark field, the CPU can compute the fermion force of the heavy fermion field. This asynchronous concurrent execution mode enhances the overall performance by $\sim 5\%$.

4. GPU optimizations

The Nvidia CUDA programming guide (see <http://developer.download.nvidia.com>) has provided general guidelines to optimize an application code for CUDA. Here we outline some useful techniques in tuning our CUDA code for CG (with mixed precision) of the (optimal) domain-wall quark matrix. First, we have one thread take care of all computations at each (s, x, y, z) , with a

loop going over all t (even/odd). The threads are divided into a one-dimensional grid with two-dimensional blocks. Each block has 64 threads, with N_s threads in the X-direction and $64/N_s$ threads in the Y-direction. The number of blocks is equal to $N_s N_x N_y N_z / 64$. At first, a naive transcription of our C code to CUDA only yielded 15 Gflops on a single GTX280. After putting the constant matrix X^{-1} (only involving the 5-th dimensional and Dirac indices) in the constant memory space, and using the texture memory for link variables and vectors, we obtained an enhancement from 15 to 30 Gflops. Then we unrolled the short loops (involving Dirac and color indices) in our CG code. This gave a further enhancement from 30 to 50 Gflops. Next we reordered the data in the global memory such that the memory access by all threads of a half-wrap is coalesced into one memory transaction. This enhanced the performance from 50 to 80 Gflops. Then we reduced the number of local variables in our code, which gave another jump from 80 to 100 Gflops. Finally, we combined the threads by re-using forward/backward data (in t) for neighboring sites. This increased the performance from 100 to 120 Gflops. For Nvidia T10/GTX285, our CG code attains 100/140 Gflops, where the difference is mainly due to the memory bandwidth and clock frequency of the GPU. For the T10 processor with 4 GB device memory, it can only accommodate a lattice up to the size $24^3 \times 48 \times 16$. In order to simulate a larger lattice, we are porting our CG CUDA code to multi-GPUs under MPI and OpenMP.

5. Current status of production runs

The first production run in two-flavor QCD is on-going, using the Iwasaki gauge action on a set of lattices with sizes $16^3 \times (32, 10, 8, 6, 4) \times (16, 32)$ at the lattice spacing $a \sim 0.1$ fm ($\beta = 2.20$), with eight sea quark masses ranging from $m_q a = 0.01, 0.02, \dots, 0.08$. After discarding 300 trajectories for thermalization, we are accumulating 5120 trajectories in total for each sea quark mass. Using the binning method, we estimate the autocorrelation time to be around 20, based on the saturation of the jackknife error of the plaquette as a function of the bin size, as shown in Fig. 1(a). Thus we sample one configuration every 20 accepted trajectories, and we have 256 configurations for each m_q . The topological charge distribution of 66 configurations for $m_q = 0.02$ is plotted in Fig. 1(b). Currently, we are completing the statistics for each gauge ensemble, and we expect to have the first physics results from the ensemble with the lattice size $16^3 \times 32 \times 16$.

6. Preliminary results of pseudoscalar mass and decay constant

In this section, we present some preliminary results of the pseudoscalar mass and decay constant, for the gauge configurations generated with 2-flavor QCD with optimal domain-wall quarks and the Iwasaki gluon action at $\beta = 2.20$, on the $16^3 \times 32 \times 16$ lattice. The weights $\{\omega_s\}$ are fixed such that the error of the sign function $S_{opt}(H_w)$ is less than 10^{-5} . In Fig. 2, we plot m_π^2/m_q and f_π versus m_q respectively. The statistics for each quark mass is about 100-150 configurations. We fit our results to the next-to-leading order (NLO) formulas in chiral perturbation theory (ChPT) [17]

$$\frac{m_\pi^2}{m_q} = 2B \left(1 + \frac{1}{2} x \ln x \right) + c_3 x, \quad x = \frac{4B}{(4\pi f)^2} m_q, \quad (6.1)$$

$$f_\pi = f(1 - x \ln x) + c_4 x. \quad (6.2)$$

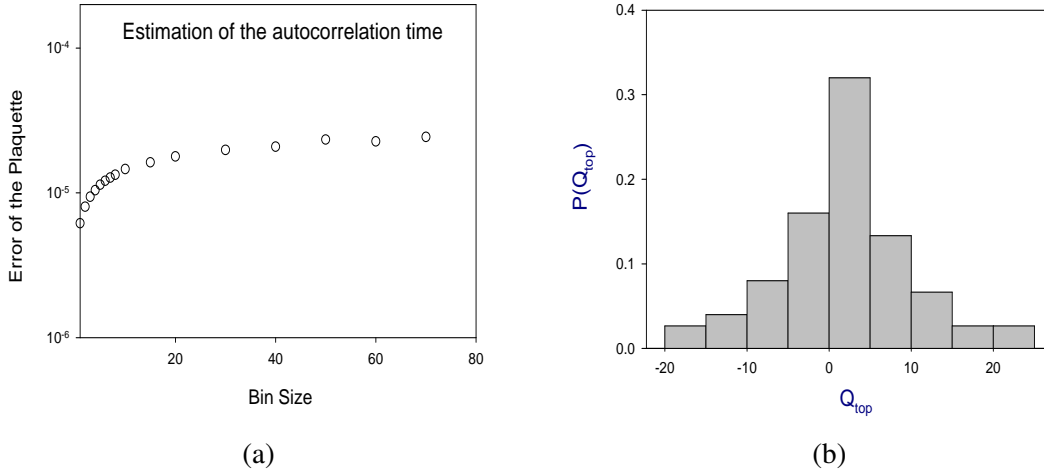


Figure 1: (a) Estimation of the autocorrelation time using the binning method. The jackknife error of the plaquette is measured from 4000 accepted trajectories in the HMC simulation of 2-flavor QCD with optimal domain-wall quarks ($m_q a = 0.02$) on the $16^3 \times 32 \times 16$ lattice. (b) The topological charge distribution of 66 configurations which are sampled every 20 accepted trajectories.




Figure 2: Preliminary results of 2-flavor QCD with optimal domain-wall quarks: (a) m_π^2/m_q , and (b) f_π . The solid lines are the fits to the ChPT at NLO.

Using the fitted value of $f a$ and the physical input $f_\pi = 131$ MeV, we obtain $a^{-1} = 1.5(1)$ GeV, a rough estimate of the lattice spacing. From the fitted value of B , we obtain the chiral condensate $\Sigma = B f^2 / 2 = [261(20)\text{MeV}]^3$. The complete analysis with the statistics of 250 configurations for each quark mass will be presented in a forthcoming paper.

7. Concluding remark

This is the first project to use a GPU cluster to perform a large-scale simulation of lattice QCD with exact chiral symmetry. Our GPU cluster is less than 1/100 in price/performance ratio and only

1/10 of the power consumption of a conventional supercomputer. Using the optimal domain-wall fermion, we can simulate unquenched lattice QCD with exact chiral symmetry, without fixing the topology. The cost of preserving chiral symmetry to a high precision is counterbalanced by the low price/performance ratio of a GPU. The production runs for 2-flavor QCD on the $16^3 \times 32 \times 16$ lattice will be completed by the end of 2009. For the (2+1)-flavor QCD on the $16^3 \times 32 \times 16$ lattice, we simulate the strange quark with our newly constructed positive-definite pseudofermion action for a single fermion flavor [18]. The thermalization of this set of gauge ensembles is now in progress.

Acknowledgments

This work is supported in part by the National Science Council (Nos. NSC96-2112-M-002-020-MY3, NSC96-2112-M-001-017-MY3, NSC98-2119-M-002-001), and NTU-CQSE (Nos. 98R0066-65, 98R0066-69). T.W.C. also thanks Nvidia Corporation for the support under Professor Partnership Program.

References

- [1] D. B. Kaplan, Phys. Lett. B **288**, 342 (1992)
- [2] H. Neuberger, Phys. Lett. B **417**, 141 (1998); R. Narayanan and H. Neuberger, Nucl. Phys. B **443**, 305 (1995)
- [3] C. Jung, PoS **LATTICE2009**, 002 (2009).
- [4] M. A. Clark, PoS **LATTICE2009**, 003 (2009).
- [5] G. I. Egri, Z. Fodor, C. Hoelbling, S. D. Katz, D. Nogradi and K. K. Szabo, Comput. Phys. Commun. **177**, 631 (2007)
- [6] H. Neuberger, Int. J. Mod. Phys. C **10**, 1051 (1999)
- [7] T. W. Chiu and T. H. Hsieh, Phys. Rev. E **68**, 066704 (2003)
- [8] K. Barros, R. Babich, R. Brower, M. A. Clark and C. Rebbi, PoS **LATTICE2008**, 045 (2008)
- [9] K. I. Ishikawa, PoS **LATTICE2008**, 013 (2008)
- [10] Z. Fodor, S. D. Katz and K. K. Szabo, JHEP **0408**, 003 (2004)
- [11] C. Allton *et al.* [RBC and UKQCD Collaborations], Phys. Rev. D **76**, 014504 (2007)
- [12] H. Fukaya, S. Hashimoto, K. I. Ishikawa, T. Kaneko, H. Matsufuru, T. Onogi and N. Yamada [JLQCD Collaboration], Phys. Rev. D **74**, 094505 (2006)
- [13] T. W. Chiu, Phys. Rev. Lett. **90**, 071601 (2003)
- [14] T. Takaishi and P. de Forcrand, Phys. Rev. E **73**, 036706 (2006)
- [15] J. C. Sexton and D. H. Weingarten, Nucl. Phys. B **380**, 665 (1992).
- [16] M. Hasenbusch, Phys. Lett. B **519**, 177 (2001)
- [17] J. Gasser and H. Leutwyler, Nucl. Phys. B **250**, 465 (1985).
- [18] K. Ogawa, T.W. Chiu, T.H. Hsieh [TWQCD Collaboration], PoS **LATTICE2009**, 033 (2009).