#### S9330: Lattice QCD with Tensor Cores

# Kate Clark<sup>3</sup>, Chulwoo Jung<sup>2</sup>, Robert Mawhinney<sup>1</sup>, and Jiqun Tu<sup>1</sup>\*

<sup>1</sup>Columbia University <sup>2</sup>Brookhaven National Labratory <sup>3</sup>Nvidia Corporation

GPU Technology Conference, Silicon Valley, March 2019

\*speaker

# Some Background

What's inside a proton?



3/36

• QCD is an  $\mathrm{SU}(3)$  local gauge theory with 6 flavors of dynamic quarks.

3/36

• QCD is an  $\frac{\mathrm{SU}(3)}{\mathrm{local gauge}}$  theory with 6 flavors of dynamic quarks.

- QCD is an SU(3) local gauge theory with 6 flavors of dynamic quarks.
- Shows *confinement*: not possible to extract individual quarks from their bound states,



Figure -1.5: Inseparability of quarks and antiquarks in spite of investing ever more energy

3/36

- QCD is an SU(3) local gauge theory with 6 flavors of dynamic quarks.
- Shows *confinement*: not possible to extract individual quarks from their bound states,



Figure -1.5: Inseparability of quarks and antiquarks in spite of investing ever more energy

 Non-perturbative: need numerical simulations to calculate measurable quantities from the theory to be measured in real world laboratories.

### The Standard Model of Particles



#### The Path Integral formulation



# Path Integral: Monte Carlo

• Path Integral formulation

$$\langle \mathcal{O} \rangle = \frac{\int [d\bar{\psi}] [d\psi] [dU] \mathcal{O}[U, \bar{\psi}, \psi] e^{-S[U, \bar{\psi}, \psi]}}{\int [d\bar{\psi}] [d\psi] [dU] e^{-S[U, \bar{\psi}, \psi]}}$$

$$S[U, \bar{\psi}, \psi] = (\text{gauge action}) + \bar{\psi}D[U]\psi$$

#### Path Integral: Monte Carlo

Path Integral formulation

$$\langle \mathcal{O} \rangle = \frac{\int [d\bar{\psi}] [d\psi] [dU] \mathcal{O}[U, \bar{\psi}, \psi] e^{-S[U, \bar{\psi}, \psi]}}{\int [d\bar{\psi}] [d\psi] [dU] e^{-S[U, \bar{\psi}, \psi]}}$$

$$S[U, \bar{\psi}, \psi] = (\text{gauge action}) + \bar{\psi}D[U]\psi$$

• Monte Carlo: Draw samples of U (gauge field configurations) according to the weight  $e^{-S}$  and measure some observable  $\mathcal{O}$  on these samples.

$$\langle \mathcal{O} \rangle \simeq \frac{1}{n} \sum_{i} \mathcal{O}_{i}$$

#### From theory to simulation

discretize the theory and put it onto a 4 dimensional lattice



Quarks cause (at least) two problems. :(

1 They are anti-commutating fermion fields/numbers,

$$\psi(x)\psi(y) = -\psi(y)\psi(x),$$

how to simulate anti-commutating numbers on a traditional computer?

Quarks cause (at least) two problems. :(

1 They are anti-commutating fermion fields/numbers,

$$\psi(x)\psi(y) = -\psi(y)\psi(x),$$

how to simulate anti-commutating numbers on a traditional computer?

2 Upon discretization they break chiral symmetry (the doubler problem): this symmetry says the nature uses left and right hand equally frequently.

Solution to the first problem:

$$\int [d\bar{\psi}][d\psi]e^{-\bar{\psi}D[U]\psi} = \det D[U] = \int [d\phi^{\dagger}][d\phi]e^{-\phi^{\dagger}D[U]^{-1}\phi}$$

The anti-commuting grassmann variables  $\psi$ 's are replaced with the usual commuting complex number  $\phi$ 's, at the expense of inverting the Dirac matrix D[U]. Typically we want to solve Dirac equations of the form

$$Dx = y, x = D^{-1}y$$

#### From theory to simulation

One solution to the second problem: we use the formulation of domain wall fermion (DWF) and add a 5th dimension



• Our numerical problem is as the following: Need to solve for a large sparse matrix of size  $\sim 10^{10}$  by  $\sim 10^{10}$  that spans over several different dimensions:

- Our numerical problem is as the following: Need to solve for a large sparse matrix of size  $\sim 10^{10}$  by  $\sim 10^{10}$  that spans over several different dimensions:
- 4d-spatial dimensions ( $\sim 10^8)$  entangled with spin(4) and color(3)

- Our numerical problem is as the following: Need to solve for a large sparse matrix of size  $\sim 10^{10}$  by  $\sim 10^{10}$  that spans over several different dimensions:
- 4d-spatial dimensions (  $\sim 10^8)$  entangled with spin(4) and color(3)
- the 5th dimension ( $\sim 10^1$ ) entangled with spin(4)

#### The Matrix

• The Dirac matrix,



#### The Matrix

• The Dirac matrix,



• The 4d-spatial operation  $(D_w)$  involves only the nearest neighbor with the same 5th dimension index, i.e. only operates horizontally.

#### The Matrix

• The Dirac matrix,



- The 4d-spatial operation  $(D_w)$  involves only the nearest neighbor with the same 5th dimension index, i.e. only operates horizontally.
- The 5th dimension operations  $(M_5^{-1} \text{ and } M_{\phi})$  involves only elements with the same 4d-spatial index, i.e. only operates vertically.

13/36

• Conjugate gradient (CG): a Krylov space iterative algorithm that minimizes the residual in each iteration.

- Conjugate gradient (CG): a Krylov space iterative algorithm that minimizes the residual in each iteration.
- Each iteration applies the normal operator dslash once.

- Conjugate gradient (CG): a Krylov space iterative algorithm that minimizes the residual in each iteration.
- Each iteration applies the normal operator dslash once.
- Use CG to solve the normal operator problem instead of the original one, († = transpose + complex conjugate)

$$Dx = y \to D^{\dagger}Dx = D^{\dagger}y.$$

13/36

- Conjugate gradient (CG): a Krylov space iterative algorithm that minimizes the residual in each iteration.
- Each iteration applies the normal operator dslash once.
- Use CG to solve the normal operator problem instead of the original one, († = transpose + complex conjugate)

$$Dx = y \to D^{\dagger}Dx = D^{\dagger}y.$$

• The convergence rate decreases as the condition number  $\kappa$  (=max eigenvalue/min eigenvalue) of the matrix  $D^{\dagger}D$  increases.

- Conjugate gradient (CG): a Krylov space iterative algorithm that minimizes the residual in each iteration.
- Each iteration applies the normal operator dslash once.
- Use CG to solve the normal operator problem instead of the original one, († = transpose + complex conjugate)

$$Dx = y \rightarrow D^{\dagger}Dx = D^{\dagger}y.$$

- The convergence rate decreases as the condition number  $\kappa$  (=max eigenvalue/min eigenvalue) of the matrix  $D^{\dagger}D$  increases.
- For our matrix  $\kappa \sim 10^8 {\rm , \ i.e.}$  extremely ill-conditioned.

- .
- Conjugate gradient (CG): a Krylov space iterative algorithm that minimizes the residual in each iteration.
- Each iteration applies the normal operator dslash once.
- Use CG to solve the normal operator problem instead of the original one, († = transpose + complex conjugate)

$$Dx = y \to D^{\dagger}Dx = D^{\dagger}y.$$

- The convergence rate decreases as the condition number  $\kappa$  (=max eigenvalue/min eigenvalue) of the matrix  $D^{\dagger}D$  increases.
- For our matrix  $\kappa \sim 10^8 {\rm , \ i.e.}$  extremely ill-conditioned.
- This matrix inversion accounts for over 90% of the simulation time.

• In the measurement phase of the Monte Carlo for one gauge field configuration typically a large number of Dirac equations with the same Dirac matrix but different RHS are solved.

$$D^{\dagger}Dx = D^{\dagger}y.$$

 In the measurement phase of the Monte Carlo for one gauge field configuration typically a large number of Dirac equations with the same Dirac matrix but different RHS are solved.

$$D^{\dagger}Dx = D^{\dagger}y.$$

 Several eigen-space methods, including the implicitly restarted Lanczos algorithm with Chebyshev polynomial [Y. Saad, SIAM J. Numer. Anal. 17, 687 (1980)], have been developed. Low-lying eigenvectors of the matrix are computed to deflate the inversions, the cost of which is amortized by the large number of RHS.

 In the measurement phase of the Monte Carlo for one gauge field configuration typically a large number of Dirac equations with the same Dirac matrix but different RHS are solved.

$$D^{\dagger}Dx = D^{\dagger}y.$$

- Several eigen-space methods, including the implicitly restarted Lanczos algorithm with Chebyshev polynomial [Y. Saad, SIAM J. Numer. Anal. 17, 687 (1980)], have been developed. Low-lying eigenvectors of the matrix are computed to deflate the inversions, the cost of which is amortized by the large number of RHS.
- Sheer size of the eigenvectors poses challenge to both the generation (not enough memory) and the storage (not enough disk space, easily reaching peta-byte scale).

#### Compression

A compression algorithm, which exploits the physical correlation between the eigenvecors, has been developed and applied. For the gauge configurations generated in this work the memory and disk space usage is expected to be reduced by a factor of 30. [arXiV:1710.06884]

The eigen-space methods does NOT work in the sample generation phase of lattice QCD since during this phase for the same Dirac matrix we only solve very few (typically one) Dirac equations.

Now finally let's put this problem onto the GPUs

 Code in this work is developed under the framework of QUDA (https://github.com/lattice/quda), a library for performing calculations in lattice QCD on GPUs.

Now finally let's put this problem onto the GPUs

- Code in this work is developed under the framework of QUDA (https://github.com/lattice/quda), a library for performing calculations in lattice QCD on GPUs.
- Matrix: Even though a sparse matrix, calculating the entries on the fly from the gauge field beats directly loading the whole matrix.

Now finally let's put this problem onto the GPUs

- Code in this work is developed under the framework of QUDA (https://github.com/lattice/quda), a library for performing calculations in lattice QCD on GPUs.
- Matrix: Even though a sparse matrix, calculating the entries on the fly from the gauge field beats directly loading the whole matrix.
- Vector: storage strategy that meets GPUs' global memory coalescing requirement.

Now finally let's put this problem onto the GPUs

- Code in this work is developed under the framework of QUDA (https://github.com/lattice/quda), a library for performing calculations in lattice QCD on GPUs.
- Matrix: Even though a sparse matrix, calculating the entries on the fly from the gauge field beats directly loading the whole matrix.
- Vector: storage strategy that meets GPUs' global memory coalescing requirement.
- Solver: Lower precisions (half/single) are used to improve the performance and higher precisions (double) are used to systematically correct the inaccuracy.

# **SUMMIT**

#### Move Over, China: U.S. Is Again Home to World's Speediest Supercomputer

June 8, 2018



# SUMMIT at ORNL

• The SUMMIT machine at Oak Ridge National Laboratory (ORNL) is currently the fastest supercomputer in the world. 6 Tesla V100 GPUs per node with a total of 4608 nodes. Our proposed jobs run on 1024 nodes (6096 GPUs) at a time.

	we have	for one iter. we need
memory bandwidth	$850~{\rm GB/s}$	69 GB
network bandwidth	$8.3~{ m GB/s}$	21 GB
network/memory	0.010	0.304

# SUMMIT at ORNL

- The SUMMIT machine at Oak Ridge National Laboratory (ORNL) is currently the fastest supercomputer in the world. 6 Tesla V100 GPUs per node with a total of 4608 nodes. Our proposed jobs run on 1024 nodes (6096 GPUs) at a time.
- Distribute our 4d-spatial lattice to the GPUs. Consequence: need to communicate between the nodes; network bandwidth largely determines the scaling.

	we have	for one iter. we need
memory bandwidth	$850~{\rm GB/s}$	69 GB
network bandwidth	$8.3~{ m GB/s}$	21 GB
network/memory	0.010	0.304

• Precondition the original matrix A to reduce the condition number  $\kappa.$ 

with preconditioning	we have	for one iter. we need
memory bandwidth	$850~{\rm GB/s}$	$69{ imes}12.2~{ m GB}$
network bandwidth	$8.3~{ m GB/s}$	21 GB
network/memory	0.010	0.025

- Precondition the original matrix A to reduce the condition number  $\kappa.$
- Solve  $(AP^{-1})(Px) = y$  instead of Ax = y. Choose some preconditioner P that requires no communication.

with preconditioning	we have	for one iter. we need
memory bandwidth	$850~{\rm GB/s}$	$69{ imes}12.2~{ m GB}$
network bandwidth	$8.3~\mathrm{GB/s}$	21 GB
network/memory	0.010	0.025

- Precondition the original matrix A to reduce the condition number  $\kappa.$
- Solve  $(AP^{-1})(Px) = y$  instead of Ax = y. Choose some preconditioner P that requires no communication.
- Consequence: Less iterations, less communication needed; Needs more local computation on each GPU.

with preconditioning	we have	for one iter. we need
memory bandwidth	$850~{\rm GB/s}$	$69{ imes}12.2~{ m GB}$
network bandwidth	$8.3~{ m GB/s}$	21 GB
network/memory	0.010	0.025

Choose the predicontioner with the *correct* boundary condition. [arXiv:1811.08488]



The preconditioner inversion  $P^{-1}$  does not need to be solved to arbitrarily high precision for the algorithm to work.



• Without preconditioning network bandwidth is the bottleneck due to large communication demand.

- Without preconditioning network bandwidth is the bottleneck due to large communication demand.
- With preconditioning the balance is shifted. In order to achieve a speed up in terms of time to solution we need to do better on utilizing the memory bandwidth and compute flops available.

- Without preconditioning network bandwidth is the bottleneck due to large communication demand.
- With preconditioning the balance is shifted. In order to achieve a speed up in terms of time to solution we need to do better on utilizing the memory bandwidth and compute flops available.
  - Memory bandwidth: kernel fusion;

- Without preconditioning network bandwidth is the bottleneck due to large communication demand.
- With preconditioning the balance is shifted. In order to achieve a speed up in terms of time to solution we need to do better on utilizing the memory bandwidth and compute flops available.
  - Memory bandwidth: kernel fusion;
  - Compute flops: tensor core for the 5th dimension operations.

The 4d-volume ( $\sim 10^8$ ) is much larger than the size of the 5th dimension ( $\sim 10^1$ ). Fusing the 4d-spatial operations with the 5th dimension operations right after reduces global memory traffic.

$$\left[1-\kappa_b^2\underbrace{M_\phi^{\dagger}D_w^{\dagger}}_{\text{fuse}}\underbrace{M_5^{-\dagger}M_\phi^{\dagger}D_w^{\dagger}}_{\text{fuse}}\underbrace{M_5^{-\dagger}\right]\left[1-\kappa_b^2M_5^{-1}D_w}_{\text{fuse}}\underbrace{M_\phi M_5^{-1}D_w}_{\text{fuse}}M_\phi\right]$$

This goes as the following:

• For each block apply the 4d-spatial operation  $D_w$ . Here we rely on the GPU cache system to achieve data reuse. The theoretical limit is 8-way (nearest neighbors in 4 directions) and we are getting around 50-60% of that.

This goes as the following:

- For each block apply the 4d-spatial operation  $D_w$ . Here we rely on the GPU cache system to achieve data reuse. The theoretical limit is 8-way (nearest neighbors in 4 directions) and we are getting around 50-60% of that.
- The results are stored in the shared memory.

This goes as the following:

- For each block apply the 4d-spatial operation  $D_w$ . Here we rely on the GPU cache system to achieve data reuse. The theoretical limit is 8-way (nearest neighbors in 4 directions) and we are getting around 50-60% of that.
- The results are stored in the shared memory.
- the 5th dimension operators are applied on the shared memory. We are not at the mercy of the cache system here: 100% data reuse.

#### **Kernel Fusion**



# Kernel Fusion: Shared Memory

27/36



Distribute the 4d-spatial indices into blocks (blockDim.x)



#### **Tensor Cores**

"Tensor cores provide a huge boost to convolutions and matrix operations."



Half precision? No problem, we are already using half precision.

#### **Tensor Cores**

"Tensor cores provide a huge boost to convolutions and matrix operations."



- Half precision? No problem, we are already using half precision.
- Memory bandwidth? No problem, the numbers are already in shared memroy.

# Memory Layout for Tensor Cores



29/36

 Due to the way the wmma functions are called it is not necessary to have two separate pieces of memory for the input and output vector. We load from and store into the same shared memory location in each block.

#### #include <mma.h>

#### using namespace nvcuda::wmma;

\_\_global\_\_ void wmma\_ker(half \*a, half \*b, float \*c) {
 // Declare the fragments
 fragment<matrix\_a, 16, 16, 16, half, col\_major> a\_frag;
 fragment<matrix\_b, 16, 16, 16, half, row\_major> b\_frag;
 fragment<accumulator, 16, 16, 16, float> c\_frag;

// Initialize the output to zero
fill\_fragment(c\_frag, 0.0f);

// Load the inputs load\_matrix\_sync(a\_frag, a, 16); load\_matrix\_sync(b\_frag, b, 16);

// Perform the matrix multiplication
mma\_sync(c\_frag, a\_frag, b\_frag, c\_frag);

// Store the output
store\_matrix\_sync(c, c\_frag, 16, wmma::mem\_row\_major);
}

# Further optimizations: Template Everything 31/36

#### • Template everything: maximize compile time optimization.





•  $12 \times 16 = 192$  threads in the block; 192/32 = 6 virtual warps.



- $12 \times 16 = 192$  threads in the block; 192/32 = 6 virtual warps.
- The numbers in each tile shows the warps that need to load it.



- $12 \times 16 = 192$  threads in the block; 192/32 = 6 virtual warps.
- The numbers in each tile shows the warps that need to load it.
- We can reload the matrix when ever needed, or preload them at the beginning.

 "To reload, or not to reload, that is the question." – William Shakespeare

- Preloading naturally saves the loading time. Reloading saves register usage: need only one wmma:fragment object.
- We use QUDA's intrinsic autotuning feature to decide which strategy wins: just try both and pick the faster one.

# Further optimizations: Spill or Not?

• Since register usage is a problem, we also tune register usage using \_\_launch\_bounds\_\_(maxThreadsPerBlock, minBlocksPerMultiprocessor).

# Further optimizations: Spill or Not?

- Since register usage is a problem, we also tune register usage using \_\_launch\_bounds\_\_(maxThreadsPerBlock, minBlocksPerMultiprocessor).
- Kernel variables could spill into device memory, potentially leads to latency.

# Further optimizations: Spill or Not?

- Since register usage is a problem, we also tune register usage using \_\_launch\_bounds\_\_(maxThreadsPerBlock, minBlocksPerMultiprocessor).
- Kernel variables could spill into device memory, potentially leads to latency.
- But why worry? It won't do any harm since we can also use autotuning to do the dirty work for us ... :)

## Results



# Conclusion

• With a carefully chosen preconditioner we are able to shift the balance between network bandwidth, GPU memory bandwidth and compute flops when solving a large sparse ill-conditioned matrix on a large scale supercomputer. The scaling is improved and we are able to achieve a speed up.

# Conclusion

- With a carefully chosen preconditioner we are able to shift the balance between network bandwidth, GPU memory bandwidth and compute flops when solving a large sparse ill-conditioned matrix on a large scale supercomputer. The scaling is improved and we are able to achieve a speed up.
- Tensor cores are used, possibly for the first time, to speed up the 5th dimension operations of domain wall fermion in lattice QCD simulations.