

SCALING LATTICE QCD ON MODERN GPU SYSTEMS

Mathias Wagner

AGENDA

Lattice Quantum Chromodynamics

GPU refresh

lower precision

multigrid

more nodes

Summary

LATTICE QUANTUM CHROMODYNAMICS

Theory is highly non-linear \Rightarrow cannot solve directly

Must resort to numerical methods to make predictions

Lattice QCD

Discretize spacetime \Rightarrow 4-d dimensional lattice of size $L_x \times L_y \times L_z \times L_t$

Finite spacetime \Rightarrow periodic boundary conditions

 $PDEs \Rightarrow$ finite difference equations

Consumer of 10-20% of public supercomputer cycles

Traditionally highly optimized on every HPC platform for the past 30 years



STEPS IN AN LQCD CALCULATION

1. Generate an ensemble of gluon field configurations "gauge generation"

Produced in sequence, with hundreds needed per ensemble Strong scaling required with 100-1000 TFLOPS sustained for several months 50-90% of the runtime is in the linear solver

O(1) solve per linear system Target 16⁴ per GPU

2. "Analyze" the configurations

Can be farmed out, assuming ~10 TFLOPS per job Task parallelism means that clusters reign supreme here 80-99% of the runtime is in the linear solver Many solves per system, e.g., O(10⁶) Target 24⁴-32⁴ per GPU $D_{ij}^{\alpha\beta}(x,y;U)\psi_{j}^{\beta}(y) = \eta_{i}^{\alpha}(x)$ or Ax = b

Simulation Cost ~ a⁻⁶ V^{5/4}



QUDA

- "QCD on CUDA" http://lattice.github.com/quda (open source, BSD license)
- Effort started at Boston University in 2008, now in wide use as the GPU backend for BQCD, Chroma, CPS, MILC, TIFR, tmLQCD, etc.

• Provides:

Various solvers for all major fermionic discretizations, with multi-GPU support Additional performance-critical routines needed for gauge-field generation

• Maximize performance

- Exploit physical symmetries to minimize memory traffic
- Mixed-precision methods
- Autotuning for high performance on all CUDA-capable architectures
- Domain-decomposed (Schwarz) preconditioners for strong scaling
- Eigenvector and deflated solvers (Lanczos, EigCG, GMRES-DR)
- Multi-source solvers
- Multigrid solvers for optimal convergence
- A research tool for how to reach the exascale

QUDA CONTRIBUTORS

10 years - lots of contributors

Ron Babich (NVIDIA) Simone Bacchio (Cyprus) Michael Baldhauf (Regensburg) Kip Barros (LANL) Rich Brower (Boston University) Nuno Cardoso (NCSA) Kate Clark (NVIDIA) Michael Cheng (Boston University) Carleton DeTar (Utah University) Justin Foley (Utah -> NIH) Joel Giedt (Rensselaer Polytechnic Institute) Arjun Gambhir (William and Mary) Steve Gottlieb (Indiana University) Kyriakos Hadjiyiannakou (Cyprus)

Dean Howarth (BU) Bálint Joó (Jlab) Hyung-Jin Kim (BNL -> Samsung) Bartek Kostrzewa (Bonn) Claudio Rebbi (Boston University) Hauke Sandmeyer (Bielefeld) Guochun Shi (NCSA -> Google) Mario Schröck (INFN) Alexei Strelchenko (FNAL) Jigun Tu (Columbia) Alejandro Vaquero (Utah University) Mathias Wagner (NVIDIA) Evan Weinberg (NVIDIA) Frank Winter (Jlab)

MAPPING THE DIRAC OPERATOR TO CUDA

Finite difference operator in LQCD is known as Dslash Assign a single space-time point to each thread

V = XYZT threads, e.g., V = 24^4 => 3.3×10^6 threads

Looping over direction each thread must

Load the neighboring spinor (24 numbers x8)

Load the color matrix connecting the sites (18 numbers x8)

Do the computation

Save the result (24 numbers)

Each thread has (Wilson Dslash) 0.92 naive arithmetic intensity

QUDA reduces memory traffic

Exact SU(3) matrix compression (18 => 12 or 8 real numbers)

Use 16-bit fixed-point representation with mixed-precision solver



NVIDIA DGX-1



8 V100 GPUs (16/32 GB) Hypercube-Mesh NVLink 4 EDR IB



NVIDIA POWERS WORLD'S FASTEST SUPERCOMPUTER

Summit Becomes First System To Scale The 100 Petaflops Milestone





27,648 Volta Tensor Core GPUs

TESLA V100 32GB

5,120 CUDA cores 640 NEW Tensor cores 7.8 FP64 TFLOPS | 15.7 FP32 TFLOPS | 125 Tensor TFLOPS 20MB SM RF | 16MB Cache 32GB HBM2 @ 900GB/s | 300GB/s NVLink



ACCELERATED COMPUTING 10X PERFORMANCE & 5X ENERGY EFFICIENCY FOR HPC

CPU Optimized for Serial Tasks



CPU Strengths

Optimized for

- Very large main memory
- Very fast clock speeds
- Latency optimized via large caches
- Small number of threads can run very quickly

CPU Weaknesses

- Relatively low memory bandwidth
- Cache misses very costly
- Low performance/watt

ACCELERATED COMPUTING 10X PERFORMANCE & 5X ENERGY EFFICIENCY FOR HPC

GPU Strengths

- High bandwidth main memory
- Significantly more compute resources
- Latency tolerant via parallelism
- High throughput
- High performance/watt

GPU Weaknesses

- Relatively low memory capacity
- Low per-thread performance

GPU Accelerator

Optimized for Parallel Tasks



SPEED V. THROUGHPUT

Speed

Throughput





Which is better depends on your needs...

*Images from Wikimedia Commons via Creative Commons

VOLTA

MAXWELL

KEPLER

PASCAL

TESLA

FERMI

QUDA STENCIL PERFORMANCE

GPUs now drive 10x faster than they used to ...

SCALING

A single GPU is not enough

Lattice QCD needs to be scaled - in one of many ways:

Multiple meanings

Same problem size, more nodes, more GPUs

Same problem, next generation GPUs

Same problem, fewer bits to solve (mixed precision)

Multigrid - strong scaling within the same run (not discussed here)

To tame strong scaling we have to understand the limiters

MIXED PRECISION

LINEAR SOLVERS

QUDA supports a wide range of linear solvers CG, BiCGstab, GCR, Multi-shift solvers, etc.

Condition number inversely proportional to mass Light (realistic) masses are highly singular Naive Krylov solvers suffer from critical slowing down at decreasing mass

Entire solver algorithm must run on GPUs Time-critical kernel is the stencil application Also require BLAS level-1 type operations while $(|\mathbf{r}_{k}| \geq \varepsilon)$ { $\beta_{k} = (\mathbf{r}_{k}, \mathbf{r}_{k})/(\mathbf{r}_{k-1}, \mathbf{r}_{k-1})$ $\mathbf{p}_{k+1} = \mathbf{r}_{k} - \beta_{k}\mathbf{p}_{k}$ $\mathbf{q}_{k+1} = \mathbf{A} \mathbf{p}_{k+1}$ $\alpha = (\mathbf{r}_{k}, \mathbf{r}_{k})/(\mathbf{p}_{k+1}, \mathbf{q}_{k+1})$ $\mathbf{r}_{k+1} = \mathbf{r}_{k} - \alpha \mathbf{q}_{k+1}$ $\mathbf{x}_{k+1} = \mathbf{x}_{k} + \alpha \mathbf{p}_{k+1}$ k = k+1}

MIXED-PRECISION CG

MIXED-PRECISION CG

Configuration provided by HotQCD collaboration (Mukherjee *et al*)

MIXED-PRECISION DEFLATION

V=48³x12, HISQ operator, physical light quarks, tol 10⁻¹⁰, 2xV100

🔄 CG double 📕 CG double-single 📃 CG double-half 📕 CG double-quarter

EIGENSOLVERS

Multiple workflows require repeated solution with different RHS with the same matrix

Multigrid not amenable to all linear operators

Eigenvector deflation is a robust alternative applicable to all operators Deflate out low modes from linear operator to accelerate the solver Cost of eigensolver is amortized if we solve enough RHS Aside: also use deflation to accelerate multigrid

Memory overheads can be limiting factor

May require storage of 1000s of vectors, ideally in fast memory

DEFLATION STABILIZES LOW PRECISION

V=48³x12, HISQ operator, physical light quarks, tol 10⁻¹⁰

16000

double-single-single

DEFLATION STABILIZES LOW PRECISION

V=48³x12, HISQ operator, physical light quarks, tol 10⁻¹⁰

16000

Solver Iterations

double-single-single

double-half-single

Tri-precision solver 12000 Outer - Inner -Eigenvector 8000 4000 Out of memory 0 32 128 0 16 64 256 512 1024 Number of eigenvectors

ονισια

Configuration provided by HotQCD collaboration (Mukherjee *et al*)

MIXED-PRECISION DEFLATION

V=48³x12, HISQ operator, physical light quarks, tol 10⁻¹⁰, 2xV100

DEFLATED ADAPTIVE MG AT THE PHYSICAL POINT

in collaboration with Dean Howarth (BU)

Multigrid shifts the lowest eigenvalues to the coarse grids

some Dirac operators (staggered / twisted mass) end up with pathological coarse-grid spectrum

For Twisted-clover operator, solution has been to add a fictitious heavy twist to the coarse operator to improve its condition number at the cost of decreased multigrid efficiency [Alexandrou et al]

Instead we deflate the coarse grid operator recovering optimal MG convergence and a 3x speedup over "mu scaling"

DEFLATED ADAPTIVE MG AT THE PHYSICAL POINT

SCALING

MULTI-GPU SCALING

HPC nodes continue to get denser

Much effort to improve strong scaling in QUDA

Key technologies employed

Peer-to-peer communication - skip MPI and directly utilize DMA copies GPU Direct RDMA (GDR) - GPU <-> NIC <-> GPU without traversing CPU memory Node topology awareness Autotuning for Dslash communication policy

Framework for Exascale communication models - NVSHMEM (OpenSHMEM for GPUs)

MULTI-GPU BUILDING BLOCKS

Halo packing Kernel

Interior Kernel

Halo communication

Halo update Kernel

BENCHMARKING TESTBED

NVIDIA Prometheus Cluster

DGX-1 nodes

8x V100 GPUs connected through NVLink4x EDR for inter-node communicationOptimal placement of GPUs and NIC

Balanced GPU / IB configuration

WHAT IS LIMITING STRONG SCALING

current state of the art using CUDA IPC (P2P) over NVLink

DGX-1,16⁴ local volume, half precision, 1x2x2x2 partitioning

WHAT IS LIMITING STRONG SCALING

Significant API overhead

REDUCING API OVERHEADS

Packing kernel writes to remote GPU using CUDA IPC

NVSHMEM

GPU centric communication

Implementation of OpenSHMEM¹, a Partitioned Global Address Space (PGAS) library NVSHMEM features

Symmetric memory allocations in device memory

Communication API calls on CPU (standard and stream-ordered)

Kernel-side communication (API and LD/ST) between GPUs

NVLink and PCIe support (intranode)

InfiniBand support (internode)

Interoperability with MPI and OpenSHMEM libraries

currently in early access

NVSHMEM DSLASH

*shmem0b.0.nvvp 🖾														-
	8.485 ms	3898.49 ms	3898.495 ms	3898.5 ms	3898.505 ms	3898.51 ms	3898.5 <mark>52.644 µs</mark> 3898	8.52 ms	3898.525 ms	3898.53 ms	3898.535 ms	3898.54 ms	3898.545 ms	3898.55
Process "dslash_testdsla							52.00							
Thread 3235466048							53 μs							
- Runtime API			cudaLa	unchKernel	cuda	aLaunchKernel	cudaLaunchKernel		cudaLau	nchKernel		cudaLa	unchKernel	
L Driver API														
Thread 2761561856														
Driver API														
Profiling Overhead														
0] Tesla V100-SXM2-16GB														
Context 1 (CUDA)														
🗆 🍸 MemCpy (HtoD)														
– 🍸 MemCpy (DtoH)								D		1	The second	11.1.		
🗆 🍸 MemCpy (DtoD)			Pac	king ke	ernel			Barr	her ker	nel	Fused	Halo		
– 🍸 MemCpy (PtoP)														
+ Compute	sonGPU <sh< td=""><td>ort void</td><td>quda::packShmemi void quda::wilso</td><td><pre>Kernel<bool=0, int:<="" ngpu<short,="" pre=""></bool=0,></pre></td><td>int=0, QudaPCTyp =4, int=3, int=1, b</td><td>e<mark>void qu</mark> oool=0, bool=0, e</td><td>ıda::syncneighborhood_k qud</td><td>ernel<bo< td=""><td>ol=1>(quda::barr</td><td>ierArg, Io</td><td>void quda::w</td><td>/ilsonGPU<sh< td=""><td>void quda::pac void qu</td><td>kShmemKer Ida::wilsonG</td></sh<></td></bo<></td></sh<>	ort void	quda::packShmemi void quda::wilso	<pre>Kernel<bool=0, int:<="" ngpu<short,="" pre=""></bool=0,></pre>	int=0, QudaPCTyp =4, int=3, int=1, b	e <mark>void qu</mark> oool=0, bool=0, e	ıda::syncneighborhood_k qud	ernel <bo< td=""><td>ol=1>(quda::barr</td><td>ierArg, Io</td><td>void quda::w</td><td>/ilsonGPU<sh< td=""><td>void quda::pac void qu</td><td>kShmemKer Ida::wilsonG</td></sh<></td></bo<>	ol=1>(quda::barr	ierArg, Io	void quda::w	/ilsonGPU <sh< td=""><td>void quda::pac void qu</td><td>kShmemKer Ida::wilsonG</td></sh<>	void quda::pac void qu	kShmemKer Ida::wilsonG
Streams				Interi	or kern	ما								
Default				intern	or kern	et								
- Stream 21		void	quda::packShmemI	Kernel <bool=0,< td=""><td>int=0, QudaPCTyp</td><td>e void qu</td><td>uda::syncneighborhood_k</td><td>ernel < bo</td><td>ol=1>(quda::barr</td><td>ierArg, lo</td><td></td><td></td><td>void quda::pac</td><td>kShmemKer</td></bool=0,<>	int=0, QudaPCTyp	e void qu	uda::syncneighborhood_k	ernel < bo	ol=1>(quda::barr	ierArg, lo			void quda::pac	kShmemKer
- Stream 24														
🗆 Stream 26														
- Stream 28														
└ Stream 29	sonGPU <sh< td=""><td>ort</td><td>void quda::wilso</td><td>nGPU<short, int:<="" td=""><td>=4, int=3, int=1, b</td><td>oool=0, bool=0, o</td><td>qud</td><td></td><td></td><td></td><td>void quda::w</td><td>/ilsonGPU<sh< td=""><td>void qu</td><td>ıda::wilsonG</td></sh<></td></short,></td></sh<>	ort	void quda::wilso	nGPU <short, int:<="" td=""><td>=4, int=3, int=1, b</td><td>oool=0, bool=0, o</td><td>qud</td><td></td><td></td><td></td><td>void quda::w</td><td>/ilsonGPU<sh< td=""><td>void qu</td><td>ıda::wilsonG</td></sh<></td></short,>	=4, int=3, int=1, b	oool=0, bool=0, o	qud				void quda::w	/ilsonGPU <sh< td=""><td>void qu</td><td>ıda::wilsonG</td></sh<>	void qu	ıda::wilsonG
- Stream 30														

DGX-1,16⁴ local volume, half precision, 1x2x2x2 partioning

NVSHMEM + FUSING KERNELS

no extra packing and barrier kernels needed

	548.45 ms	3548.455 ms	3548.46 ms	3548.465 ms	3548.47 <mark>35.903 µs</mark>	3548.475 ms	3548.48 ms	3548.485 ms	3548.49 ms
Process "dslash_testdsla									
🖃 Thread 316765952					36 us				
Driver API					p				
Thread 958622528									
Runtime API	cudaLaunchKern	el	cu	daLaunchKernel	cudal	LaunchKernel		cudaLaunchKernel	
Driver API									
Profiling Overhead									
😑 [0] Tesla V100–SXM2–16GB									
Context 1 (CUDA)									
– 🍸 MemCpy (HtoD)									
– 🍸 MemCpy (DtoH)		Interio	r + Pack +	Flag korno			Parriar	Eurod Halo	
🗆 🍸 MemCpy (DtoD)		mene	i i lack i	i lug kerne	L		Darrier	ruseu naio	
– 🍸 MemCpy (PtoP)									
Compute	t=1, bool=0	void quda::wilsonGPU	<short, int="1</td"><td>, bool=0, bool=0, quda::Kern</td><td>elType, quda::WilsonArg<shor< td=""><td>rt, int=3, Quda</td><td>void quda::wilsonGPU<</td><td>short, int=4, int=3, int=1, bool=0,</td><td>void quda</td></shor<></td></short,>	, bool=0, bool=0, quda::Kern	elType, quda::WilsonArg <shor< td=""><td>rt, int=3, Quda</td><td>void quda::wilsonGPU<</td><td>short, int=4, int=3, int=1, bool=0,</td><td>void quda</td></shor<>	rt, int=3, Quda	void quda::wilsonGPU<	short, int=4, int=3, int=1, bool=0,	void quda
🗆 🍸 40.5% void quda		void quda::wilsonGPU	<short, int="1</td"><td>, bool=0, bool=0, quda::Kern</td><td>elType, quda::WilsonArg<shor< td=""><td>rt, int=3, Quda</td><td></td><td></td><td>void quda</td></shor<></td></short,>	, bool=0, bool=0, quda::Kern	elType, quda::WilsonArg <shor< td=""><td>rt, int=3, Quda</td><td></td><td></td><td>void quda</td></shor<>	rt, int=3, Quda			void quda
∟ 🍸 37.2% barrier_all									
∟ 🍸 19.9% void quda	t=1, bool=0						void quda::wilsonGPU<	short, int=4, int=3, int=1, bool=0,	
∟ 🍸 1.1% void quda:									
⊢ 🍸 0.3% void quda::									
∟ 🍸 0.3% void quda::									
∟ 🍸 0.2% void quda::									
∟ 🍸 0.1% void quda::									
∟ 🍸 0.1% void quda:									
– 🍸 0.3% memset (0)									
Streams									
Default									

DGX-1,16⁴ local volume, half precision, 1x2x2x2 partioning

LATENCY OPTIMIZATIONS Different strategies implemented

DGX-2: FULL NON-BLOCKING BANDWIDTH

2.4 TB/s bisection bandwidth

DGX-2 STRONG SCALING

Global Volume 32⁴, Wilson-Dslash, half precision

DGX-2 STRONG SCALING

Global Volume 32⁴, Wilson-Dslash, half precision

DGX-2 STRONG SCALING

Global Volume 32⁴, Wilson-Dslash, half precision

MULTI-NODE SCALING

DGX SuperPOD (DGX2-H: 16 V100 (32GB), 8 EDR IB)

MULTI-NODE SCALING

DGX SuperPOD (DGX2-H: 16 V100 (32GB), 8 EDR IB)

NVSHMEM OUTLOOK

	548.45 ms	3548.455 ms	3548.46 ms	3548.465 ms	3548.47 <mark>35.903 µs</mark>	3548.475 ms	3548.48 ms	3548.485 ms	3548.49 ms
Process "dslash_testdsla	1								
🖃 Thread 316765952									
Driver API									
Thread 958622528									
Runtime API	cudaLaunchKernel		cud	aLaunchKernel	cudal	aunchKernel		cudaLaunchKernel	
Driver API									
Profiling Overhead									
📃 [0] Tesla V100–SXM2–16GB									
🚍 Context 1 (CUDA)									
🗆 🍸 MemCpy (HtoD)									
🗕 🍸 MemCpy (DtoH)									
🗆 🍸 MemCpy (DtoD)									
– 🍸 MemCpy (PtoP)									
🖃 Compute	t=1, bool=0	void quda::wilsonGPU<	short, int=4, int=3, int=1,	bool=0, bool=0, quda::Kerno	elType, quda::WilsonArg <shor< td=""><td>t, int=3, Quda</td><td>void quda::wilsonGPU<shor< td=""><td>t, int=4, int=3, int=1, bool=0,</td><td>void quda</td></shor<></td></shor<>	t, int=3, Quda	void quda::wilsonGPU <shor< td=""><td>t, int=4, int=3, int=1, bool=0,</td><td>void quda</td></shor<>	t, int=4, int=3, int=1, bool=0,	void quda
占 🍸 40.5% void quda		void quda::wilsonGPU<	short, int=4, int=3, int=1,	bool=0, bool=0, quda::Kern	elType, quda::WilsonArg <shor< td=""><td>t, int=3, Quda</td><td></td><td></td><td>void quda</td></shor<>	t, int=3, Quda			void quda
- 🍸 37.2% barrier_all									
∟ 🍸 19.9% void quda	t=1, bool=0						void quda::wilsonGPU <shor< td=""><td>t, int=4, int=3, int=1, bool=0,</td><td></td></shor<>	t, int=4, int=3, int=1, bool=0,	
∟ 🍸 1.1% void quda:									
∟ 🍸 0.3% void quda::									
∟ 🍸 0.3% void quda::									
∟ 🍸 0.2% void quda::									
∟ 🍸 0.1% void quda::									
∟ 🍸 0.1% void quda:									
L 🍸 0.3% memset (0)									
Streams									
Default									

NVSHMEM OUTLOOK

	548.45 ms	3548.455 ms	3548.46 ms	3548.465 ms	3548.47 <mark>35.903 µs</mark>	3548.475 ms	3548.48 ms	3548.485 ms	3548.49 ms
Process "dslash_testdsla									
🖃 Thread 316765952									
Driver API									
Thread 958622528									
Runtime API	cudaLaunchKernel		cu	daLaunchKernel	cuda	LaunchKernel		cudaLaunchKernel	
Driver API									
Profiling Overhead									
📃 [0] Tesla V100–SXM2–16GB									
Context 1 (CUDA)									
– 🍸 MemCpy (HtoD)									
🗆 🍸 MemCpy (DtoH)									
🗆 🍸 MemCpy (DtoD)									
– 🍸 MemCpy (PtoP)									
🖃 Compute	t=1, bool=0								void quda
占 🍸 40.5% void quda		IOne kerne	el to rule i	them all !					void quda
∟ 🍸 37.2% barrier_all		C		والمحالية المحالية			ويتعامله فجار ومتر		
∟ 🍸 19.9% void quda	t=1, bool=0	Communic	ation is na	andled in the	e kernel and	latencies a	are nidden.		
∟ 🍸 1.1% void quda:									
⊢ 🍸 0.3% void quda::									
∟ 🍸 0.3% void quda::									
∟ 🍸 0.2% void quda::									
∟ 🍸 0.1% void quda::									
∟ 🍸 0.1% void quda:									
L 🍸 0.3% memset (0)									
Streams									
Default									

IN THE WILD

CHROMA HMC MULTIGRID

HMC typically dominated by solving the Dirac equation, but Few solves per linear system Can be bound by heavy solves (c.f. Hasenbusch mass preconditioning)

Multigrid setup must run at speed of light

Reuse and evolve multigrid setup where possible

- Use the same null space for all masses (setup run on lightest mass)
- Evolve null space vectors as the gauge field evolves (Lüscher 2007)
- Update null space when the preconditioner degrades too much on lightest mass

CHROMA HMC-MG ON SUMMIT

From Titan running 2016 code to Summit running 2019 code we see >82x speedup in HMC throughput

Multiplicative speedup coming from machine and algorithm innovation Highly optimized multigrid for gauge field evolution Force gradient integrator

Chroma

Bálint Joó

DPP-JI7

Boram Yoon -Frank Winter

NODE PERFORMANCE OVER TIME

Multiplicative speedup through software and hardware

Time to solution is measured time to solution for solving the Wilson operator against a random source on a 24x24x24x64 lattice, β =5.5, M_{π} = 416 MeV. One node is defined to be 3 GPUs

QUDA - LATTICE QCD ON GPUS

Widely used for Lattice QCD applications on GPUs

State of the art solvers using mixed precision

Multigrid Deflation Block-Krylov solver

All components for gauge field evolution

Portable high-performance kernels through auto-tuning and careful optimization Tuned Multi-GPU scaling

GPU centric communication with NVSHMEM takes CPU limitations out Multiplicative speedup from hardware and software: more science

