

anton antigene accurs The first descent from 27.1 Institute accurs an

e anten el gette there to object Lafer Manufato (1011) (111)

CHREATLY CONTRACT RANK

Application of multigrid solvers in lattice QCD simulations at physical quark masses

02. October 2023

Seminar

RIKEN Kobe

Jacob Finkenrath



Table of Content	Wuppertal2010 -DDalphaAMGAndreas FrommerMatthias RottmannKarst	2015 PhD
1. Physics* Short motivation	Bjoern LederFrancesco KnechliA. Frommer et. al., SIAM J.Sci.Comput. 36 (2014)A. Frommer et. al., arXiv:1307.6101	
2. HardwareSetting the stage		The Cyprus Institute2015 - 2022Adaptation to twisted massAcceleration of HMC
 3. Multigrid solvers Introduction 	Cor	Alexandrou Simone Bacchio Giannis Koutsou Shuhei Yamamoto C. Alexandrou et al., PRD 94 (2016) 11, 114509
 Acceleration of the HMC Coarse-grid improvements 	Wuppertal2022 - nowCoarse grid improvementsMulti-level trace estimation	S. Yamamoto et al., PoS LATTICE2021 (2022) 536
OR 5269	Gustavo Ramírez Jesus Espinoza-Valverde Jose Jimenez Henning Leemhuis J. Espinoza-Valverde et al.,CPC 292 (2023) 1088 A. Frommer et al., <i>SIAM J.Sci.Comput.</i> 44 (2022) 4, A25	Artur Strobel 369 36-A2556

Standard model



Understanding the building blocks of our world







Exascale is reached with Frontier: a lot of available compute time

However, scaling up is non-trivial due to increase of complexity on the single but also multi-node



Machines



Challenges

Best = 1 , worse >> 1

 10^{3}





- Node level performance is increasing
- Single Nodes become a power house usually equipped with several accelerators (higher complexity to optimize)



- On node bandwidth does not grow as fast as FP
- Limiting factor for lattice QCD kernels
- Scalable algorithms are becoming more complex

Flops over interconnect



- Interconnect bandwidth is way behind FP node performance
- Extrem scaling is very likely limited by interconnect

P. Boyle, Lattice 2023, Plenary, Advances in algorithms for solvers and gauge generation

Introduction

- Aller



Multigrid solvers

From BlueGene to IntelXeon

- higher performance on single node





Linear equation



Linear equation

Computational challenge solving:

$$Dx = b$$
 $D \in \mathbb{C}^{n imes n}$ and $x, b \in \mathbb{C}^n$

 * 4-dim matrix stencil with $n\sim 10^9$





- matrix vector stencil with nearest neighbour connections
 Iterative linear solver approach requires matrix-vector product
- Matrix vector product bandwidth bounded:

Flops : Bytes = 1.3:1.1





Benchmark



Performance of matrix vector product

UEABS Benchmark-kernels on PRACE Tier 0 systems (PRACE 4/5IP)

Application of lattice codes:

- QUDA for Nvidia GPUs
- QPhiX for Intel CPUs
- Grid for Arm

Computation cost dominated by matrix vector application



Computational performance is bounded by available bandwidth
 Towards 16 nodes, Nvidia GPU deviates from perfect scaling



Krylov Subspace Methods

Iterative solver: Krylov space

$$x \in \mathscr{K}_n(D,b) = \left\{ b, Db, D^2b, \dots, D^{n-1}b \right\}$$

And solve subsystem

Cost for Conjugate gradient solver:

$$cost\approx V\cdot\left(\frac{b}{\mu}+a\right)\approx V\frac{b}{\mu}$$

with $b/a \sim 0.04$



Iteration proportional to condition number

only depend on smallest mode

iteration count increases drastically at physical point

Computational kernel

• Simple, mainly $D \cdot x_i$

Increase convergence by preconditioning





QUDA mixCG: (Juwels Booster)

- 64^3x128: 16 A100 GPUs
- 80^3x160: 32 A100 GPUs
- 96^3x196: 64 A100 GPUs

SAP

Idea: Preconditioning IR and UV modesError can be split:

$$\epsilon = \epsilon_{low} + \epsilon_{high}$$

* A good smoother eliminates ϵ_{high} high effectively

• Multigrid solver can lead to convergence independent of $\kappa(A)$

A good smoother: apply block inverse in alternating procedure Precondition iterate in Krylov space via:

$$\begin{array}{c} x_i \to (1-KD)x_i + Kb \\ & \quad \text{With} \quad K = D_w^{-1} + D_b^{-1} - D_b^{-1}D_{bw}D_w^{-1} \\ \text{And error:} \quad \\ \epsilon_{SAP} = 1 - KD = (1-D_b^{-1}D)(1-D_w^{-1}D) \\ \end{array}$$





Assign at least two block per node (or core) no (inter-node) communication is needed while inverting blocks Smoother simple to implement acts on the large eigenvalues

M. Luscher, Comput.Phys.Commun. 156 (2004) 209-220



 S^1O

Algebraic MG



Algebraic multigrid: non-geometric

Two-level AMG \clubsuit Petrov-Galerkin projection $D_c = RDP$

A ALLAN

✤ Coarse-grid correction:

$$\psi \leftarrow \psi + P D_c^{-1} R r$$
 with $r = \eta - D \psi$

Error propagation

$$\epsilon = I - PD_c^{-1}RD$$

Capture sufficient low modes within P

Extension to Multilevel AMG

Extension is possible: apply decomposition to the second level

Algorithm	Two-level V-cycle			
Input : ψ , η ,	ν			
Output : ψ				
$r \leftarrow \eta - D\psi$				
$\psi \leftarrow \psi + P D_c^{-1} R r$				
$r \leftarrow \eta - D\psi$	·)			
$\psi \leftarrow \psi + PD$ $r \leftarrow \eta - D\psi$ $\psi \leftarrow \psi + M^{(i)}$	$c^{-1}Rr$			

-



Aggregation





Aggregation blocks:

$$\mathcal{A}_i = \mathcal{L}_{j(i)} imes \mathcal{W}_i$$

Keel Ar

$$\{\mathcal{A}_{i,\tau}: j=1,\ldots,n_{\mathcal{L}_c}, \tau=0,1\}$$

Where

$$\mathcal{A}_{i,0} = \mathcal{L}_j \times \{0,1\} \times \mathcal{C}$$
$$\mathcal{A}_{i,1} = \mathcal{L}_j \times \{2,3\} \times \mathcal{C}$$

this implies

$$\Gamma_5 P = P \Gamma_5^c$$



Test vector Aggregation:







Building up the coarse grid space

Test vectors: use approximate eigenvectors in the construction of P
An ideal approach: Take eigenspace of small eigenvalues
done in setup phase

Idea: use multigrid hierarchy to approximated D^{-1}

Update multigrid hierarchy as the test vectors improve

 $v_{i+1} \leftarrow "D^{-1}"v_i$

M. Luscher, JHEP 12 (2007) 011, JHEP 07 (2007) 081



-



Setup



A. Frommer et. al., arXiv:1307.6101



Multigrid solver is almost independent of condition number

DDalphaAMG

Domain Decomposed adaptive algebraic Multigrid solver

 Up to two magnitude faster at the physical point compare to BiCGstab or CG
 Comes with an additional overhead, coarse grid projection has to be build

A. Frommer et al., SIAM J.Sci.Comput. 36 (2014)



HMC and MG

and the second



Acceleration of HMC

From IntelXeon to GPUs





-

Simulation in QCD







FOR 5269



Hybrid Monte Carlo with IR-UV preconditioning

Target: sample configuration distributed with the Boltzmann weight

 $P(U) \propto \det[D^{\dagger}D(U)] \exp\{-\beta S_g(U)\}$

Idea: use IR+UV filtering via:
Hasenbusch mass - preconditioning
Nested integrators

$$\det[D^{\dagger}D] = \det\left[\frac{D^{\dagger}D}{D^{\dagger}D + \mu^{2}}\right] \cdot \det[D^{\dagger}D + \mu^{2}]$$





Fermion force requires inversions

✤ integrate all terms with mu<0.01 with MG</p>

Re-write the determinants:

Challenges in HMC

Enabling MG solver in Molecular Dynamics



Requires single solves:

$$(DD^{\dagger} + \mu^2)^{-1} = \gamma_5 (D - i\gamma_5 \mu)^{-1} \gamma_5 (D + i\gamma_5 \mu)^{-1}$$

Cost efficient if one can

- re-use setup phase, adequate update procedure during trajectory
- re-use coarse levels for all operators
 - for different flavors, up, down, strange and charm quarks
 - ✤ with different shifts

by ensuring stability of method during molecular dynamics

Outcome of DDalphaAMG adaptation in HMC:

- very stable
- need one update iteration for setup every light quark force

(20 - 50 updates during 1 MDU)



Challenges in HMC



requirements 5x

parallelization :

hard limit for minimal

multigrid increases memory

L=64 ~ 16 Nodes

L=96 ~ 80 Nodes

Memory:

*

*

*

*



Scalability:

- Limited by the size of the coarsest grid
- on SuperMUC-NG window scales with

 $\propto V^{\frac{3}{4}}$





Reversibility:

- reused of multigrid subspace can impact reversibility
- using a higher solver precision for force computation:

factor 10 higher precision for MG



Speedup

ETMC Simulation effort on SuperMUC



Pioneering Multigrid solver (work with S. Bacchio)

- Nf =2 physical point Simulations
 - two ensembles at larger box sizes

C. Alexandrou et al., PRD 94 (2016) 11, 114509

- Inclusion of heavy quark doublet
 - ongoing: four ensembles at different lattice spacings

C. Alexandrou et al., CPC. 236 (2019), 51-64

- Higher order integrators
 - fourth order integrator

D. Shcherbakov, CCP 21 (2017) 4, 1141-1153



However, for exa-scale eraadaptation to GPU machines are (likely) necessary



20

Adaptation to GPUs

Work with

Bartosz Kostrzewa	Marco Ga	arofalo	Simone Romiti
Sime	ne Bacchio	Feren	Pittler



 Idea: Use tmLQCDs linkage to QUDA
 QUDA comes with highly optimized solvers for GPUs (includes now also AMD and Intel kernels)

Similar to MG-adaptation

outsource computation intensive parts to the GPU

- Start with Inversions
 - enable multigrid solver
 - ✤ enable multi-mass-shift CG
- ♦ (on-going) Extend to force calculation
 - Pure gauge part available
 - Currently fermion force is under work



MG on GPU is even further limited (due to memory and scalability)

CG solver more efficient due to utilisation of mix-precision



GPU utilisation

FOR 5269

B. Kostrzewą, arXiv:2212.06635

Breakdown of time spend in monomials Marconi 100 (16 Nodes with 4 V100)



~ 36% associated with MG solver Other parts better scalable

Example of CPU dominant ndcloverrat1 derivative





Example of GPU dominant





22

HMC on GPUs

FOR 5269







HMC accelerated with GPU, enables MD times below 2 hours
Autotuner for MG parameters turned out to be very useful
Allows in principle larger lattices, e.g. with L=192
Enables us to utilize top machines in EU and US
e.g. Lumi in Finnland

23

Strong scaling

The second se



Strong scaling of MG

Overcoming network bounds



24

-

BERGISCHE UNIVERSITÄT WUPPERTAL

Coarsest level improvements

Three level method

SuperMUC-NG scaling



- ✤ Although the coarsest level is solved to 10^-1,
 - dominating in the strong scaling limit
- Coarse grid becomes too small, communication dominates

Juwels-Booster scaling





Coarse grid

Improvements

Need for coarse grid improvements



Coarsest level improvements

The Cyprus Institute

S. Yamamoto, PoS LATTICE2021 (2022) 536

Simone Bacchio

Shuhei Yamamoto J.F.

and the second

Enable multi-right-hand sides

- increase computational load compare to bandwidth
- Block-Krylov methods
- Simple vectorisations

Bergische Universitaet Wuppertal

J. Espinoza-Valverde et al., arXiv:2205.09104

Gustavo Ramírez

Jesus Espinoza-Valverde

Introduced four new methods for coarsest-level

- Block-diagonal preconditioned: no communication involved
- polynomial preconditioned: exchanges dot products for matrix-vector multiplications, which improves scaling
- Deflation and recycling (via GCRO-DR): reduces the iteration count, which reduces the number of dot products, which in turns improves scalability
- pipelined (does not improve)

Multi - rhs

DDalphaAMG multiple RHS



Implementation:

- definition of a new data structure, index on vectors runs the fastest
 - ✤ re-written all low level kernels
- using auto-vectorisation to enhance portability
- On Skylake chips this shifts vectorisation from 128 bit to 256 bit
- Coarse grid computation speed up by a factor ~1.8





Combination with Block-Krylov solvers possible

- Iteration count is indeed reduced
- Additional overhead to high so far
 Additional overhead orthogonalise

Re-ordering and orthogonalisation

FOR 5269

Multi - rhs



Multiple rhs increase scalability

comes with an additional memory overhead

✤ up to a factor 5 on Intel Skylake machines

✤ result so far only obtained on Intel and AMD CPUs

S. Yamamoto, PoS LATTICE2021 (2022) 536





DDalphaAMG - Strong scaling

V = 64x32x32x

0

Ο

Ο

V = 160x80x80x80

V = 192x96x96x96

✨

О

0

O single rhs - native

ideal scaling

10²

rel. speedup for a single rhs

 10^{1}

0⁰ 10

10⁰

multi rhs, 4 rhs (new)

multi rhs, 8 rhs (new)



 10^{4}





Coarsest level improvements

J. Espinoza-Valverde et al., CPC 292 (2023) 108869

Algorithm #1: block-diagonal preconditioned

A state of the sta

$$Lets write \qquad D_L = \begin{pmatrix} D_{ee} & D_{eo} \\ D_{oe} & D_{oo} \end{pmatrix}$$

- * then the Schur complement is given by $D_c = D_{ee} - D_{eo} D_{oo}^{-1} D_{oe}$
- ✤ We can then focus on

$$D_c x = b$$

which can be preconditioned via:

$$D_{ee}^{-1}D_cx = D_{ee}^{-1}b$$



m_0	without BDP	with BDP
-0.3515	21	15
-0.35371847789	35	26
-0.354	40	28
-0.3545	52	37

1



29



Coarsest level improvements

Algorithm #2: Polynomial preconditioned

Idea: use $q(D_c) \approx D_c^{-1}$

Then:
$$D_c q_d(D_c) y = b$$
 and $x = q_d(D_c) y$
And $q_{\mu-1}(t) = \sum_{i=1}^{\mu} \frac{1}{\theta_i} \prod_{j=1}^{i-1} \left(1 - \frac{t}{\theta_j}\right)$ with θ_i Ritz values

Algorithm #3: GCRO-DR

- Polynomial spreads the spectrum around zero, which is beneficial if we want to do deflation: $(I-C_kC_k^H)D_c$
- ✤ by extracting k eigenvalues of the current Krylov cycle via $U = [u_1, \ldots u_k] \quad \text{and} \quad C_k = D_c U_k$
- \clubsuit include the vectors $[c_1,\ldots,c_k]$ $\;$ in the next Krylov cycle

Spectrum transformation: $D_c \rightarrow D_c q_{15}(D_c)$



Figure 3.5.(a) in ``Toward efficient polynomial preconditioning for GMRES" by Loe and Morgan



30

Coarsest level improvements

J. Espinoza-Valverde et al., arXiv:2205.09104



- On JUWELS cluster
- ✤ lattice 128x64x64x64
- Wilson-Dirac discretisation

� k = 25, d=4

- Block diagonal preconditioned
- Polynomial preconditioned
- ✤ GCRO-DR

m0 = -0.355937
 (most ill-conditioned)





Overview: Multigrid improvements



32



coarse grid deflation



openBC fully supported

Currently working on a comparison of available MG-methods

Conclusion



Conclusion

Multigrid is the method of choice for Wilson like fermions

- for physical light quarks
- ✤ accelerate HMC simulations
- Coarse grid improvements needed to extend strong scaling

Multigrid method can be also utilised in different cases:

- Trace estimation (Multi-Level Monte Carlo)
 - ♦Use coarser grid to split up trace into IR and UV dominated pieces

A. Frommer et al., SIAM J. Sci. Comput. 44 (2022) 4, A2536-A2556

Outlook:

New DDalphaAMG lib written in C++,

- ♦use to test additional methods
 - Like LU-decomposition of the coarse grid



Next step: combination with multi-level algorithms ...

1

