

Deflation

Steven Gottlieb

Questions


- To what extent can critical slowing down be slowed by deflation or reweighting?
- Caveats:
 - I don't work in this area
 - I don't know the literature well
 - I did consult with Kostas Orginos, but he is not responsible for my (mis)interpretations

Accelerated DD-HMC

M. Lüscher wrote two papers last year on deflation:

- “Local coherence and deflation of the low quark modes in lattice QCD,” JHEP **0707**, 081 (2007) [arXiv:0706.2298 [hep-lat]].
- “Deflation acceleration of lattice QCD simulations,” JHEP **0712**, 011 (2007) [arXiv:0710.5417 [hep-lat]].
- The first paper deals with domain-decomposed deflation subspaces used for propagator calculation.
- The second paper applies this idea to configuration generation.
- Improvement is substantial: more than a factor of three on lightest mass.

DD-HMC Performance Enhancement

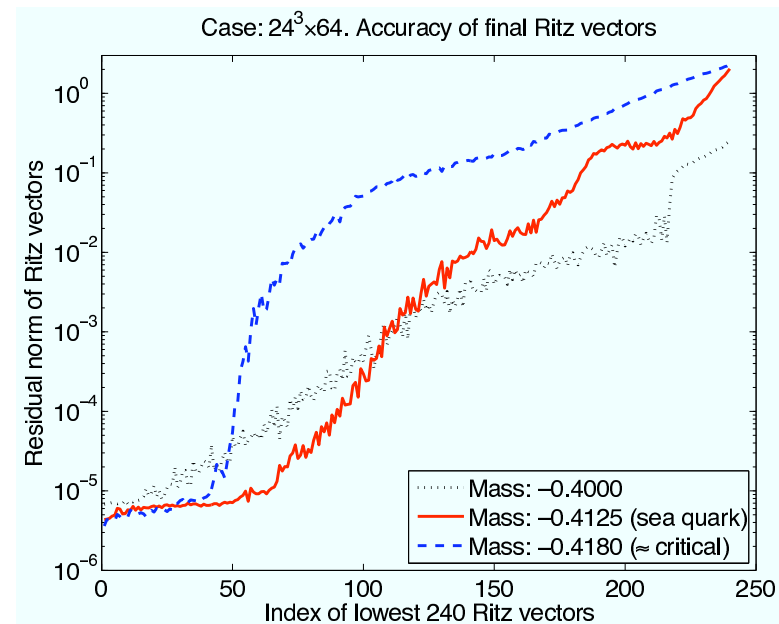
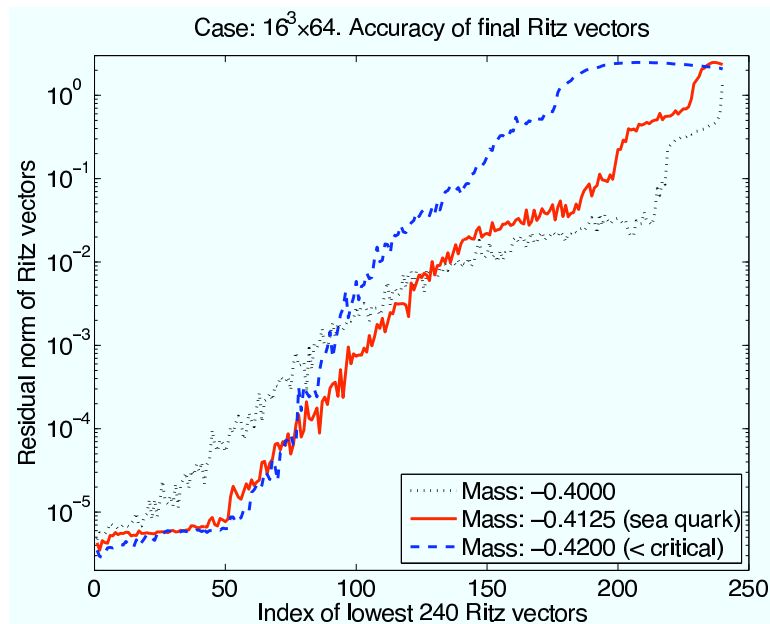


Execution time t required per molecular-dynamics trajectory of length $\tau = 0.5$ vs. the inverse of the bare current-quark mass m_{sea} given in lattice units. At the smallest quark mass considered, the time needed for two trajectories of length $\tau = 0.25$ is plotted. Lattice size is $32^3 \times 64$. Computer is 32 dual-socket single-core nodes with Infiniband network.

Deflation for Multiple RHSs


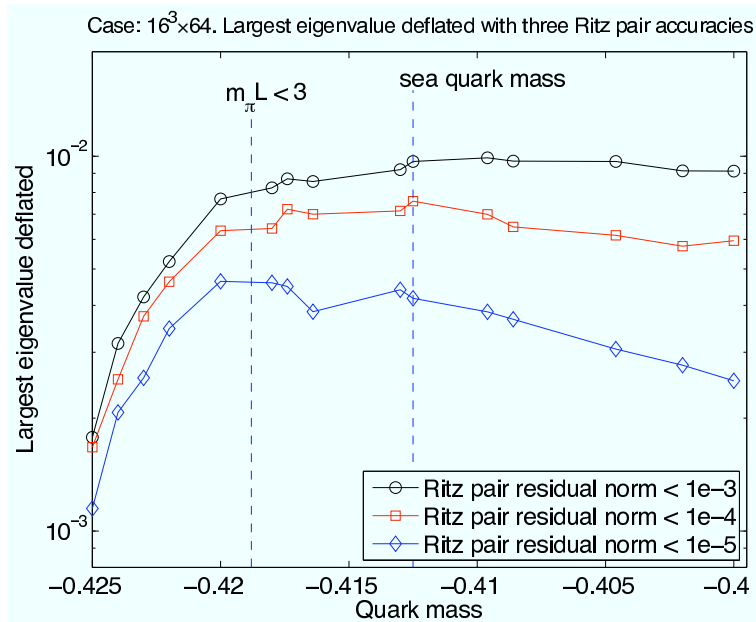
- A. Stathopoulos and K. Orginos, Computing and deflating eigenvalues while solving multiple right hand side linear systems with an application to Quantum Chromodynamics, arXiv:0707.0131 [hep-lat]
 - eigCG(n_eigenvectors, iterations_between_restarts)
 - Incremental eigCG: solve a new RHS and add another n_eigenvectors to the space of eigenvectors for deflation
 - initCG: deflation applied to initial guess
- N-N scattering lengths require 1500–3000 RHSs per configuration
- Tests on $24^s \times 64$ lattice with 256 RHSs on Cray XT4

Accuracy of Eigenvalues




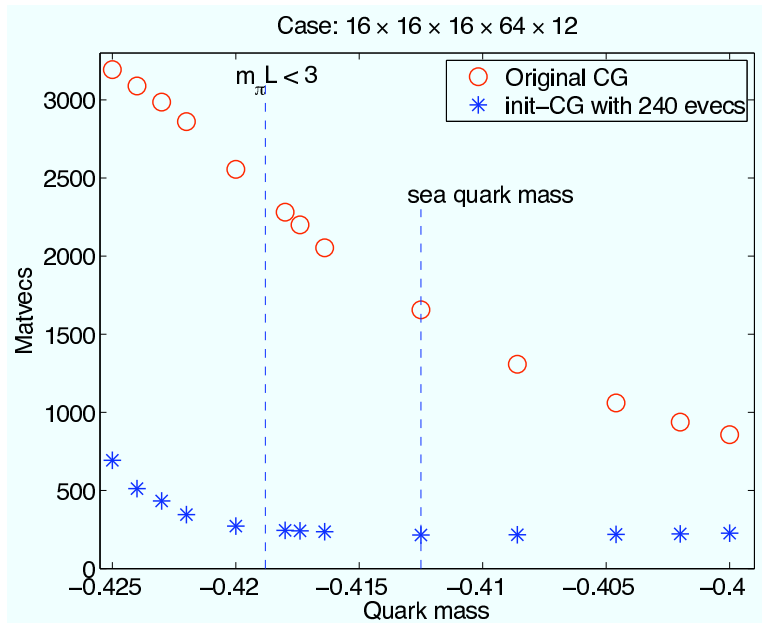
Residual norm of the 240 Ritz vectors computed at the end of Incremental eigCG on 24 right hand sides. Left graph shows results from the 3M lattice and right graph from the 10M lattice. The three curves correspond to three different quark masses (a heavy, the sea quark, and one close to the critical mass).

Largest Eigenvalue Convergence



For each quark mass and its corresponding matrix we plot the largest eigenvalue of the 240 computed by Incremental eigCG that has residual norm less than some threshold. Each of the three curves corresponds to a different threshold.

Removal of Critical Slowdown



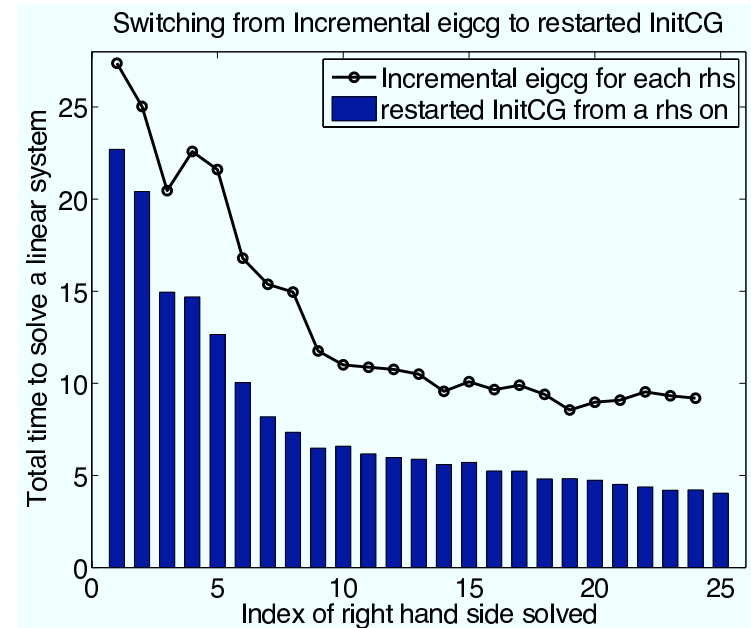
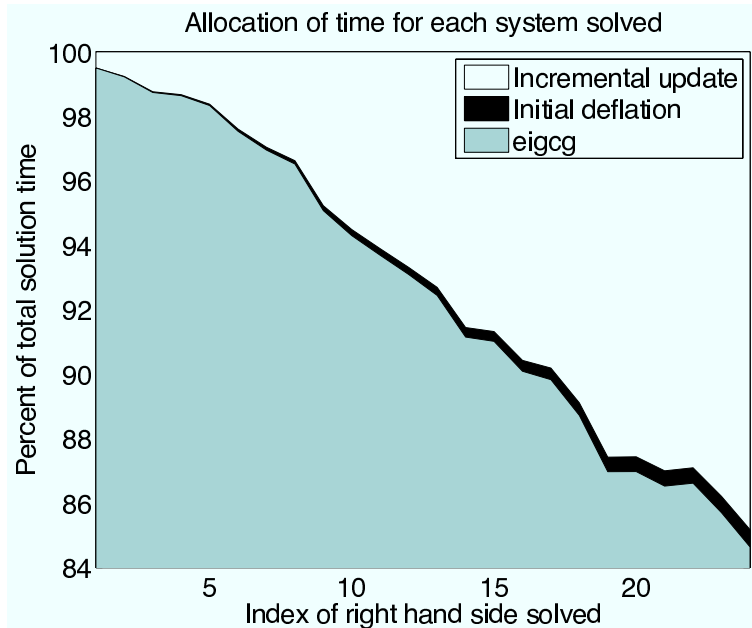
For each quark mass and its corresponding matrix we plot the average number of iterations required by init-CG to solve the rest 24 systems. For comparison the number of iterations of the non-deflated CG is reported.

Overall Speedup

Execution times in seconds for phase 1 where we apply Incremental eigCG, for phase 2 where we switch to init-CG, and for the whole application that includes also about 90 seconds of application setup time. We also show the times from using the original CG in Chroma without deflation.

Time for first 24 rhs		Time for next 232 rhs		Total application time	
Chroma CG	527.9	Chroma CG	5127.2	Chroma CG	5751.5
Incr. eigCG	323.2	init-CG	951.2	Incr. eigCG	1365.8
speedup	1.6	speedup	5.4	overall speedup	4.2

Optimizing Number of Eigenvectors



L: Percentage breakdown of execution time for the three components of Incremental eigCG: the eigCG inner iteration, the initial deflation, and the incremental update. R: The solid line shows the execution time in seconds of Incremental eigCG for the k -th right hand side, $k = 1, \dots, 24$. The bars below each point show the time it would take to solve a linear system by switching to init-CG at the k -th right hand side

Optimizing Number of Eigenvectors II

total number of RHS	6	12	24	32	48
optimal # RHS in Incremental eigCG	3	7	9	11	19


Clearly, the more right hand sides we need to solve, the more vectors it pays to solve with Incremental eigCG.

Morgan/Wilcox

Reviewed at Lattice 2007 in Regensburg: W. M. Wilcox, 'Deflation Methods in Fermion Inverters,' PoS LAT2007, 025 (2007) [arXiv:0710.1813 [hep-lat]].

	Morgan/Wilcox	Stath./Orginos	Lüscher
Basic solvers	GMRES; BiCGStab	CG	GCR
Matrix type	non-hermitian (algebraic)	hermitian (algebraic)	non-hermitian (lattice)
Preconditioning?	no	no	yes (SAP)
Simultaneous solve?	yes	yes	no
Eigenvalue use for multiple rhs's	Every cycle (GMRES-Proj) At beginning (BiCGStab)	Every cycle ($s \leq s_1$); beginning and restart ($s > s_1$)	Every outer iteration
Matrix shifting?	yes (GMRES) no (BiCGStab)	no	no
Algorithm acceleration	mild	large	small

GMRES-DR & GMRES-PROJ



- GMRES-DR, builds up an eigenvalue space and uses it for deflation on restart.
- GMRES-PROJ, uses the eigenvalue space determined by GMRES-DR to solve with a second RHS.

Additional Morgan/Wilcox references

W. Wilcox, R. B. Morgan and collaborators have a number of contributions on inflation.

- A. M. Abdel-Rehim, R. B. Morgan and W. Wilcox, “Improved Seed Methods for Symmetric Positive Definite Linear Equations with Multiple Right-hand Sides,” arXiv:0810.0330 [math-ph].
- D. Darnell, R. B. Morgan and W. Wilcox, “Deflated GMRES for Systems with Multiple Shifts and Multiple Right-Hand Linear Algebra Appl. **429**, 2415 (2008) [arXiv:0707.0502 [math-ph]].
- R. B. Morgan and W. Wilcox, “Deflated Iterative Methods for Linear Equations with Multiple Right-Hand Sides,” arXiv:0707.0505 [math-ph].