

A nested Krylov subspace method for the overlap operator

Jacques C.R. Bloch*

University of Regensburg

E-mail: jacques.bloch@physik.uni-regensburg.de

Simon Heybrock

University of Regensburg

E-mail: simon.heybrock@physik.uni-regensburg.de

We present a novel method to compute the overlap Dirac operator at zero and nonzero quark chemical potential. To approximate the sign function of large, sparse matrices, standard methods project the operator on a much smaller Krylov subspace, on which the matrix function is computed exactly. However, for large lattices this subspace can still be too large for an efficient calculation of the sign function. The idea of the new method is to nest Krylov subspace approximations by making a further projection on an even smaller subspace, which is then small enough to compute the sign function efficiently, and this without any noticeable loss of numerical accuracy. We demonstrate the efficiency of the method both on Hermitian and non-Hermitian matrices.

Supported by the DFG collaborative research center SFB/TR-55 "Hadron Physics from Lattice QCD".

The XXVII International Symposium on Lattice Field Theory

July 26-31, 2009

Peking University, Beijing, China

*Speaker.

1. The overlap operator and the sign function

In this contribution we present an improved method to compute the sign function of a large sparse matrix, as necessitated by the overlap Dirac operator in lattice QCD. More details and results can be found in Ref. [1]. The new method applies to Hermitian and non-Hermitian matrices and is therefore usable in QCD simulations at both zero and nonzero quark chemical potential.

The overlap operator introduced by Neuberger and Narayanan [2, 3] and extended to nonzero quark chemical potential μ by Bloch and Wettig [4] is defined as

$$D_{\text{ov}}(\mu) = 1 + \gamma_5 \text{sgn}(\gamma_5 D_{\text{w}}(\mu)), \quad (1.1)$$

with Wilson-Dirac operator $D_{\text{w}}(\mu)$ and Wilson mass satisfying $-2 < m_{\text{w}} < 0$. The quark chemical potential is introduced in the Wilson-Dirac operator as prescribed by Hasenfratz and Karsch [5]. The matrix $\gamma_5 D_{\text{w}}$ is Hermitian when $\mu = 0$, but becomes non-Hermitian when $\mu \neq 0$, such that the overlap operator requires the computation of the sign of a general complex matrix.

For a generic function f and a diagonalizable matrix $A = U \text{diag}(\lambda_1, \dots, \lambda_n) U^{-1} \in \mathbb{C}^{n \times n}$, with eigenvalues λ_i and eigenvector matrix $U = [u_1 | \dots | u_n]$, the matrix function $f(A)$ is given by the spectral definition

$$f(A) = U \text{diag}(f(\lambda_1), \dots, f(\lambda_n)) U^{-1}. \quad (1.2)$$

In the non-Hermitian case the eigenvalues are generically complex and the sign function is defined as $\text{sgn}(z) = \text{sgn}(\text{Re}(z))$ [6].

2. Approximate Krylov subspace solution

For physically relevant lattice sizes the spectral definition (1.2) cannot be applied directly as the diagonalization of A becomes too expensive in terms of CPU time and memory. A solution is to compute $f(A)x$, for $x \in \mathbb{C}^n$, using a Krylov subspace approximation, rather than $f(A)$, as the results of such operations are typically needed by iterative solvers for linear systems and eigensystems.

To approximate $f(A)x$ in the Krylov subspace $\mathcal{K}_k(A, x) = \text{span}(x, Ax, \dots, A^{k-1}x)$ we first construct a basis for the subspace. In the Hermitian case the Lanczos method generates an orthonormal basis V_k of $\mathcal{K}_k(A, x)$ using short recurrences. In the non-Hermitian case an orthonormal basis can be constructed using the Arnoldi algorithm. However, as the method uses long recurrences, it becomes too expensive for large Krylov subspaces. Instead, we will use the two-sided or biorthogonal Lanczos method (2sL), where the orthogonality of the basis is given up in order to recover short recurrences. Two short recurrence relations are used to construct mutually orthonormal bases V_k and W_k , i.e., $W_k^\dagger V_k = \mathbf{1}_k$, for the Krylov subspaces of A and A^\dagger , respectively.

With these bases the Ritz approximation to $f(A)x$ is computed using

$$f(A)x \approx V_k f(H_k) V_k^\dagger x = \|x\| V_k f(H_k) e_1^{(k)}, \quad (2.1)$$

where we chose $v_1 = x/\|x\|$, $e_1^{(k)}$ is the first unit vector of \mathbb{C}^k , and the $k \times k$ Ritz matrix H_k is defined as $V_k^\dagger A V_k$ and $W_k^\dagger A V_k$, for the Hermitian and non-Hermitian case, respectively. The virtue of the Krylov subspace approximation is that a good accuracy can be reached for $k \ll n$, such that

the computation of $f(A)$ is replaced by that of $f(H_k)$, which is of much smaller size. For the sign function, $\text{sgn}(H_k)$ is typically computed using the Roberts-Higham matrix iteration [7]:

$$S_{n+1} = \frac{1}{2}(S_n + S_n^{-1}) \quad \text{with} \quad S_0 = A, \quad (2.2)$$

which converges quadratically to $\text{sgn}(A)$. The combination of Eq.(2.1) with Eq. (2.2) to compute $\text{sgn}(H_k)$ will be called the non-nested method, in contrast to the nested method introduced in Sec. 3.

Some additional care has to be taken when using the approximation (2.1) for the sign function. If A has eigenvalues close to the discontinuity along the imaginary axis, the Krylov subspace would have to be taken very large to achieve a good accuracy with Eq. (2.1). This can be resolved by treating these *critical eigenvalues* exactly using deflation [8]. Assume that we have computed m such critical eigenvalues λ_i with their corresponding right and left eigenvectors r_i and l_i , then the function evaluation is rewritten as $f(A)x = R_m f(\Lambda_m)(L_m^\dagger x) + f(A)(I - R_m L_m^\dagger)x$, where $\Lambda_m = \text{diag}(\lambda_1, \dots, \lambda_m)$, $R_m = [r_1 | \dots | r_m]$ and $L_m = [l_1 | \dots | l_m]$. The first contribution to $f(A)x$ is computed exactly using the eigensolutions for the deflated eigenvalues, while the second contribution is approximated using the Krylov subspace approximation (2.1).

The non-nested approximation has been implemented and tested on configurations ranging from 4^4 to $16^3 \times 32$ lattices [8–10]. The performance of the method is illustrated for a typical 8^4 configuration in Fig. 1. As expected, the convergence is generically faster and smoother in the Hermitian case (Lanczos) than in the non-Hermitian case (2sL), see Fig. 1a. More relevant for the current study is Fig. 1b which depicts the CPU times used in both cases. The figure highlights a serious issue with the non-nested method, as a bottleneck in the numerical computation can clearly be identified, in both the Hermitian and non-Hermitian case. The difference between the full and dashed lines corresponds to the time needed to compute $\text{sgn}(H_k)$ using the iteration (2.2). Evidently, this computation takes up a substantial amount of the total CPU time, as the $\mathcal{O}(k^3)$ complexity makes the Roberts-Higham iteration overly expensive for large Krylov subspaces. In this talk we will present a new method which alleviates this problem by computing $\text{sgn}(H_k)e_1^{(k)}$ of Eq. (2.1) using an additional Krylov subspace level.




Figure 1: Accuracy versus (a) Krylov subspace size (left) and (b) CPU time (right) for a typical 8^4 configuration in the Hermitian ($\mu = 0$) and non-Hermitian ($\mu = 0.3$) cases, with deflation gap $\Delta = 0.055$. In the right panel the full lines show the total time required to compute $\text{sgn}(A)x$, while the dashed lines give the time needed to build the bases in the Krylov subspaces.

3. Nested Krylov subspace method for the sign function

The idea is to approximate the vector $\text{sgn}(H_k)e_1^{(k)}$ of dimension k in a smaller, nested Krylov subspace $\mathcal{K}_\ell(H_k, e_1^{(k)})$ of dimension ℓ using the approximation (2.1), with $A \rightarrow H_k$ and $x \rightarrow e_1^{(k)}$:

$$\text{sgn}(A)x \approx \|x\|V_k \text{sgn}(H_k)e_1^{(k)} \approx \|x\|V_k V_\ell \text{sgn}(H_\ell)e_1^{(\ell)}, \quad (3.1)$$

where $\text{sgn}(H_\ell)$ is computed using Eq. (2.2). As it stands, Eq. (3.1) does not improve the efficiency of the original method because the inner Krylov subspace $\mathcal{K}_\ell(H_k, e_1^{(k)})$ only contains information coming from the $\ell \times \ell$ upper left corner of H_k , due to the tridiagonal nature of H_k and the special source vector $e_1^{(k)}$. Therefore, the size of the inner Krylov subspace has to be chosen $\ell \approx k$ in order to achieve the full accuracy of the outer Krylov subspace, and nothing has been gained.

There is, however, a way to circumvent this problem and make the idea work. We observe that

$$\text{sgn}(A) = \text{sgn}(A + A^{-1}), \quad (3.2)$$

since the eigenvectors of both arguments are the same, and the transformation does not change the sign of the eigenvalues, i.e.,

$$\text{sgn}\left(z + \frac{1}{z}\right) = \text{sgn} \text{Re}\left(z + \frac{1}{z}\right) = \text{sgn} \text{Re}\left(z + \frac{z^*}{|z|^2}\right) = \left(1 + \frac{1}{|z|^2}\right) \text{sgn} \text{Re}(z) = \text{sgn} z, \quad (3.3)$$

for $z \in \mathbb{C}$. Based on the property (3.2) we introduce a preconditioning step $H_k \rightarrow H_k + H_k^{-1}$ in Eq. (2.1), which enables us to use the nested approximation (3.1) with inner Krylov subspace $\mathcal{K}_\ell(H_k + H_k^{-1}, e_1^{(k)})$ instead of $\mathcal{K}_\ell(H_k, e_1^{(k)})$. As we will see in Sec. 4, the nested approximation with preconditioning step works well with $\ell \ll k$, i.e., there is no noticeable loss in accuracy even after reducing the Krylov subspace size substantially. This increased efficiency can be understood by noting that the transformation $H_k \rightarrow H_k + H_k^{-1}$ improves the condition number by a factor ten, approximately. This is illustrated in Fig. 2 showing the spectra of H_k and $H_k + H_k^{-1}$ for the Hermitian case. Clearly, the preconditioning step considerably widens the gap around the origin, causing the improvement in condition number.

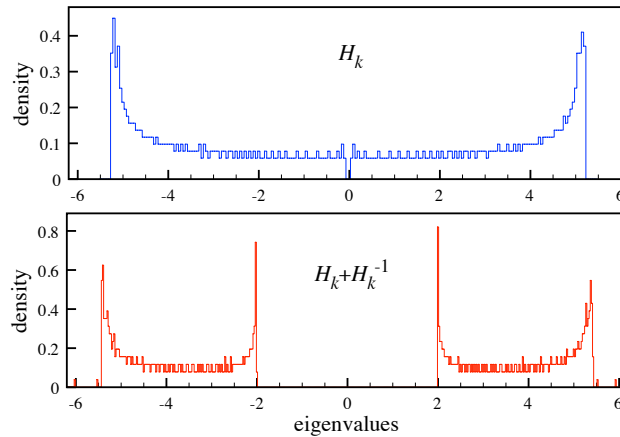


Figure 2: Eigenvalue distributions of H_k (top) and $H_k + H_k^{-1}$ (bottom) for a 6^4 lattice with $k = 1024$.




Figure 3: Convergence of the nested method for an 8^4 lattice in the Hermitian case (deflation gap $\Delta = 0.055$). The blue line shows the accuracy ε of the non-nested approximation versus the outer Krylov subspace size k . The vertical line picks a value for k , here $k = 1536$, corresponding to a desired accuracy. Finally, the red line shows the accuracy of the nested method as a function of the inner Krylov subspace size ℓ for fixed k .

4. Numerical results

In this section we show some preliminary numerical results for the nested Krylov subspace method. The gain in efficiency of the method is characterized by the smallest value of ℓ/k for which no relevant loss of accuracy occurs. This is illustrated in Fig. 3, which shows the convergence of the nested method. The crucial feature is that, for fixed k , the accuracy of the nested approximation remains optimal over a very large range in ℓ until the error eventually blows up when the inner Krylov subspace becomes too small. From the figure one observes that ℓ can be chosen about ten times smaller than k without affecting the accuracy of the approximation. Fig. 4 illustrates how this reduction translates in a gain in CPU time. For fixed k , the inner Krylov subspace size ℓ is varied and the corresponding accuracy and CPU time can be read from the figure. The large gain in CPU time achieved when reducing ℓ is due to the $\mathcal{O}(\ell^3)$ cost to compute $\text{sgn}(H_\ell)$. It is remarkable that

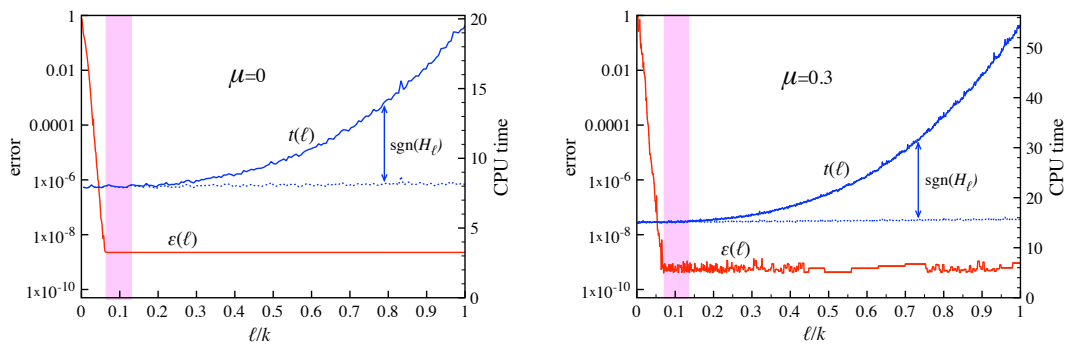


Figure 4: Gain in CPU time of the nested method for an 8^4 lattice in the Hermitian case (left) with $k = 1536$ ($\Delta = 0.055$) and the non-Hermitian case (right) for $\mu = 0.3$ with $k = 1600$ ($\Delta = 0.11$). The red line shows the accuracy $\varepsilon(\ell)$, for a fixed outer Krylov subspace size k . The total CPU time used by the method is given by the full blue line, while the horizontal dotted line represents the time needed to construct the outer Krylov subspace. The difference between both corresponds to the time needed to compute $\text{sgn}(H_\ell)$ using the Roberts-Higham iteration and is of order $\mathcal{O}(\ell^3)$. The vertical band shows the range of optimal ℓ .

there is a region of ℓ , given by the vertical band, where the accuracy is still maximal but where the computation time of $\text{sgn}(H_k)e_1^{(k)}$ is negligible. This makes the nested method extremely efficient, at least for the lattice sizes considered in these preliminary tests.

To explore the employability of the method for realistic lattice calculations it is useful to investigate how the algorithm scales with the lattice volume. This is illustrated in Fig. 5. Fitting the CPU time gives an approximate volume dependence $t \sim V^{1.2}$ for the Hermitian case and $t \sim V^{1.4}$ for the non-Hermitian case. In this context we also compare the efficiency of our new method with the rational approximation methods, which are the best methods currently on the market, see Fig. 6. For the Hermitian case the nested method performs slightly better than the Zolotarev method. However, the difference is not significant and conclusions may depend on the details of the implementation (especially since the nested method is sometimes used in double pass mode, depending on the available hardware, to avoid storage problems for large lattices [1]). For the non-Hermitian case, where we compare with the rational approximation method presented in Ref. [10], the conclusion is unambiguous and the nested method is by far better than the rational approximation. This is caused by the lesser efficiency of the rational approximation in the presence of complex eigenvalues [10].

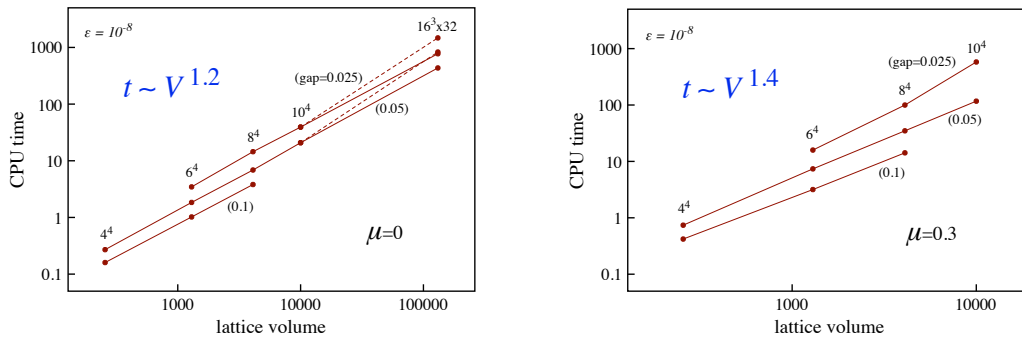


Figure 5: Scaling with matrix size: CPU time versus lattice volume for the Hermitian case (left) and non-Hermitian case with $\mu = 0.3$ (right) for three different deflation gaps: 0.025, 0.5 and 0.1. The requested accuracy is $\varepsilon = 10^{-8}$. The dashed lines represent double pass calculations [1].

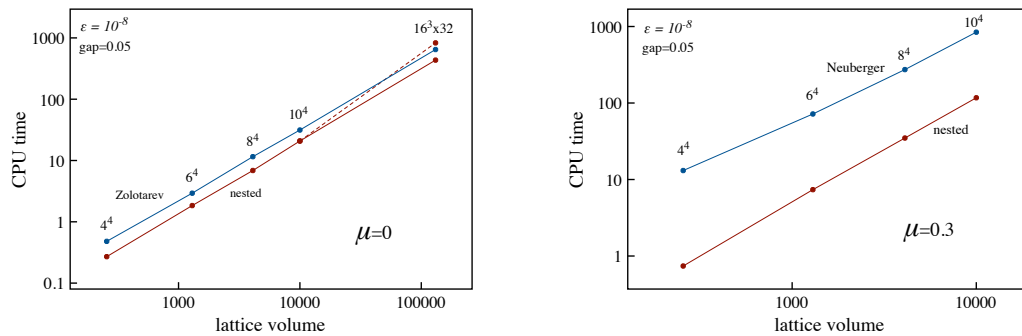


Figure 6: Comparison of the CPU time used by the nested method and rational approximation methods: For the Hermitian case (left) we compare the nested method with the Zolotarev approximation [11], for the non-Hermitian case (right) we compare with the Neuberger approximation computed with a multishift, restarted FOM method [10]. The requested accuracy is $\varepsilon = 10^{-8}$ and the deflation gap is $\Delta = 0.05$.

5. Summary and outlook

The Ritz approximation to the sign function slows down dramatically when the Krylov subspace grows large. We therefore developed an improvement based on nested Krylov subspaces, which resolves this problem and expedites the computation of the sign function for Hermitian and non-Hermitian matrices, without affecting the accuracy of the approximation. Moreover, the new method turns out to be a worthy alternative to state-of-the-art rational approximation methods. More details about the nested method can be found in Ref. [1].

Future developments will include the parallel implementation and benchmarking of the nested method, its incorporation in hybrid Monte Carlo algorithms for dynamical simulations with overlap fermions, and the investigation of the applicability of the nested method to other matrix functions.

Acknowledgments

We would like to thank Andreas Frommer and Tilo Wettig for discussions.

References

- [1] S. Heybrock and J. C. R. Bloch, *A nested Krylov subspace method to compute the sign function of a complex matrix*, in preparation.
- [2] R. Narayanan and H. Neuberger, *A construction of lattice chiral gauge theories*, *Nucl. Phys.* **B443** (1995) 305–385, [[hep-th/9411108](#)].
- [3] H. Neuberger, *Exactly massless quarks on the lattice*, *Phys. Lett.* **B417** (1998) 141–144, [[hep-lat/9707022](#)].
- [4] J. C. R. Bloch and T. Wettig, *Overlap Dirac operator at nonzero chemical potential and random matrix theory*, *Phys. Rev. Lett.* **97** (2006) 012003, [[hep-lat/0604020](#)].
- [5] P. Hasenfratz and F. Karsch, *Chemical potential on the lattice*, *Phys. Lett.* **B125** (1983) 308.
- [6] J. C. R. Bloch and T. Wettig, *Domain-wall and overlap fermions at nonzero quark chemical potential*, *Phys. Rev.* **D76** (2007) 114511, [[arXiv:0709.4630](#)].
- [7] J. Roberts, *Linear model reduction and solution of the algebraic Riccati equation by use of the sign functions*, *Internat. J. Control* **32** (1980) 677–687.
- [8] J. C. R. Bloch, A. Frommer, B. Lang, and T. Wettig, *An iterative method to compute the sign function of a non-Hermitian matrix and its application to the overlap Dirac operator at nonzero chemical potential*, *Comput. Phys. Commun.* **177** (2007) 933–943, [[arXiv:0704.3486](#)].
- [9] J. C. R. Bloch, T. Brey, and T. Wettig, *Comparing iterative methods to compute the overlap Dirac operator at nonzero chemical potential*, *PoS LATTICE2008* (2008) 027, [[arXiv:0810.4228](#)].
- [10] J. C. R. Bloch, T. Brey, A. Frommer, S. Heybrock, K. Schäfer, and T. Wettig, *Short-recurrence Krylov subspace methods for the overlap Dirac operator at nonzero chemical potential*, [arXiv:0910.1048](#).
- [11] J. van den Eshof, A. Frommer, T. Lippert, K. Schilling, and H. A. van der Vorst, *Numerical methods for the QCD overlap operator. I: Sign-function and error bounds*, *Comput. Phys. Commun.* **146** (2002) 203–224, [[hep-lat/0202025](#)].