



Bergische Universität Wuppertal

Fakultät für Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational
Mathematics (IMACM)

Preprint BUW-IMACM 21/27

Andreas Frommer and Mostafa Nasr Khalil and Gustavo
Ramirez-Hidalgo

A Multilevel Approach to Variance Reduction in the Stochastic Estimation of the Trace of a Matrix

August 23, 2021

<http://www.imacm.uni-wuppertal.de>

A MULTILEVEL APPROACH TO VARIANCE REDUCTION IN THE STOCHASTIC ESTIMATION OF THE TRACE OF A MATRIX*

ANDREAS FROMMER[†], MOSTAFA NASR KHALIL[†], AND GUSTAVO RAMIREZ-HIDALGO[†]

Abstract. The trace of a matrix function $f(A)$, most notably of the matrix inverse, can be estimated stochastically using samples $x^* f(A)x$ if the components of the random vectors x obey an appropriate probability distribution. However such a Monte-Carlo sampling suffers from the fact that the accuracy depends quadratically of the samples to use, thus making higher precision estimation very costly. In this paper we suggest and investigate a multilevel Monte-Carlo approach which uses a multigrid hierarchy to stochastically estimate the trace. This results in a substantial reduction of the variance, so that higher precision can be obtained at much less effort. We illustrate this for the trace of the inverse using three different classes of matrices.

Key words. Monte-Carlo, multilevel Monte-Carlo, trace, Hutchinson’s method, sparse matrices

AMS subject classifications. 65C05, 65F10, 65F60, 65N55

1. Introduction. We consider the situation where one is interested in the trace $\text{tr}(f(A))$ of a matrix function $f(A)$. Here, $f(A) \in \mathbb{C}^{n \times n}$ is the matrix obtained from $A \in \mathbb{C}^{n \times n}$ and the function $f : z \in D \subseteq \mathbb{C} \rightarrow f(z) \in \mathbb{C}$ in the usual operator theoretic sense; see [28], e.g. Our focus is on the inverse A^{-1} , i.e. $f(z) = z^{-1}$. Computing the trace is an important task arising in many applications. The trace of the inverse is required, for example, in the study of fractals [40], in generalized cross-validation and its applications [25, 23]. In network analysis, the Estrada index—a total centrality measure for networks—is defined as the trace of the exponential of the adjacency matrix A of a graph [16, 21] and an analogous measure is given by the trace of the resolvent $(\rho I - A)^{-1}$ [15, Section 8.1]. For Hermitian positive definite matrices A , one can compute the log-determinant $\log(\det(A))$ as the trace of the logarithm of A . The log-determinant is needed in machine learning and related fields [35, 38]. Further applications are discussed in [33, 47, 46]. A particular application which we want to prepare with this work is in lattice quantum chromodynamics (QCD), a computational approach in Theoretical Physics to simulate the interaction of the quarks as constituents of matter. Here, the trace of the inverse of the discretized Dirac operator yields the disconnected fermion loop contribution to an observable; see [42]. As simulation methods get more and more precise, these contributions become increasingly important.

It is usually unfeasible to compute the diagonal entries $f(A)_{ii}$ directly as $e_i^* f(A)e_i$, e_i the i th canonical unit vector, and then obtain the trace by summation. For example, for the inverse this would mean that we have to solve n linear systems, which is prohibitive for large values of n .

One large class of methods which aims at circumventing this cost barrier are deterministic approximation techniques. Probing methods, for example, approximate

$$(1.1) \quad \text{tr}(f(A)) \approx \sum_{i=1}^N w_i^* f(A)w_i,$$

*Submitted to the editors DATE.

Funding: This work was funded by the European Commission within the STIMULATE project.

[†]University of Wuppertal, Department of Mathematics, Gauss-Strasse 20, 42097 Wuppertal, Germany (frommer@uni-wuppertal.de, m.nasr@stimulate-ejd.eu, g.ramirez@stimulate-ejd.eu).

where the vectors w_i are carefully chosen sums of canonical unit vectors and N is not too large. Various approaches have been suggested and explored in order to keep N small while at the same time achieving good accuracy in (1.1). This includes approaches based on graph colorings; see [6, 14, 44] e.g., and the hierarchical probing techniques from [43, 30]. In order for probing to yield good results, the matrix $f(A)$ should expose a decay of the moduli of its entries when we move away from the diagonal, since the sizes of the entries farther away from the diagonal determine the accuracy of the approximation. Recent theoretical results in this direction were given in [18]. Lanczos techniques represent another deterministic approximation approach and are investigated [7, 10, 31], e.g. Without giving details let us just mention that in order to improve their accuracy, deterministic approximation techniques can be combined with the stochastic techniques to be presented in the sequel; see [46], e.g. In this paper we deal with the other big class of methods which aim at breaking the cost barrier using *stochastic estimation*. In principle, they work for any matrix and, for example, do not require a decay away from the diagonal. Our goal is to develop a multilevel Monte-Carlo method to estimate $\text{tr}(f(A))$ stochastically. Our approach can be regarded as a variance reduction technique applied to the classical stochastic ‘‘Hutchinson’’ estimator [29]

$$(1.2) \quad \text{tr}(f(A)) \approx \frac{1}{N} \sum_{n=1}^N x^{(n)} f(A) x^{(n)},$$

where the components of the random vectors $x^{(n)}$ obey an appropriate probability distribution. The variance of the estimator (1.2) decreases only like $\frac{1}{N}$, which makes the method too costly when higher precisions are to be achieved. The multilevel approach aims at curing this by working with representations of A at different levels. On the higher numbered levels, evaluating $f(A)$ becomes increasingly cheap, while on the lower levels, which are more costly to evaluate, the variance is small.

This paper is organized as follows: In Section 2 we recall the general framework of multilevel Monte-Carlo estimators. In Section 3 we discuss Hutchinson’s method for stochastically estimating the trace before turning to our new multilevel approach in Section 4. This section also contains a comparison to known approaches based on deflation as a motivation of why the new multilevel method should provide additional efficiency. Finally, several numerical results are presented in Section 5.

2. Multilevel Monte-Carlo. We discuss the basics of the multilevel Monte-Carlo approach as a variance reduction technique. We place ourselves in a general setting, thereby closely following [20].

Assume that we are given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with sample space Ω , sigma-algebra $\mathcal{F} \subseteq \Omega$ and probability measure $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$. For a given random variable $f : \Omega \rightarrow \mathbb{C}$, the *standard Monte-Carlo* approach estimates its expected value $\mathbb{E}[f]$ as the arithmetic mean

$$(2.1) \quad \mathbb{E}[f] \approx \frac{1}{N} \sum_{n=1}^N f(\omega^{(n)}),$$

where the $\omega^{(i)}$ are independent events coming from $(\Omega, \mathcal{F}, \mathbb{P})$. The variance of this estimator is $\frac{1}{N} \mathbb{V}[f]$, so the root mean square deviation has order $\mathcal{O}(N^{-1/2})$. This indicates that the number N of events has to increase quadratically with the accuracy required which is why, typically, higher accuracies require very high computational effort in this type of Monte-Carlo estimation.

The idea of *multilevel Monte-Carlo* is to split the random variable f as a sum

$$(2.2) \quad f = \sum_{\ell=1}^L g_\ell,$$

where the random variables $g_\ell : \Omega \rightarrow \mathbb{C}$ are regarded as contributions “at level ℓ ” to f . This gives

$$\mathbb{E}[f] = \sum_{\ell=1}^L \mathbb{E}[g_\ell],$$

and an unbiased estimator for $\mathbb{E}[f]$ is obtained as

$$\mathbb{E}[f] \approx \sum_{\ell=1}^L \frac{1}{N_\ell} \sum_{n=1}^{N_\ell} g_\ell(\omega^{n,\ell}),$$

where the $\omega^{(n,\ell)}$ denote the independent events on each level. The variance of this estimator is

$$\sum_{\ell=1}^L \frac{1}{N_\ell} \mathbb{V}[g_\ell].$$

The idea is that we are able to find a multilevel decomposition of the form (2.2) in which the cost C_ℓ to evaluate g_ℓ is low when the variance $V_\ell := \mathbb{V}[g_\ell]$ is high and vice versa. As is explained in [20], the solution to the minimization problem which minimizes the total cost subject to achieving a given target variance ϵ^2

$$\sum_{\ell=1}^L N_\ell C_\ell \rightarrow \min! \quad \text{s.t.} \quad \sum_{\ell=1}^L \frac{1}{N_\ell} V_\ell = \epsilon^2$$

gives $N_\ell = \mu \sqrt{V_\ell / C_\ell}$. Here, the Lagrangian multiplier μ satisfies $\mu = \epsilon^{-2} \sum_{\ell=1}^L \sqrt{V_\ell / C_\ell}$,

and the corresponding minimal total cost is

$$C = \epsilon^{-2} \left(\sum_{\ell=1}^L \sqrt{V_\ell C_\ell} \right)^2.$$

The typical situation is that the contributions g_ℓ on level ℓ are given as differences $f_\ell - f_{\ell+1}$ of approximations f_ℓ to f on the various levels, i.e. we have

$$(2.3) \quad f = \sum_{\ell=1}^{L-1} \underbrace{(f_\ell - f_{\ell+1})}_{=g_\ell} + \underbrace{f_L}_{=g_L} \quad \text{with } f_1 = f.$$

If we assume that the cost \hat{C}_ℓ to evaluate f_ℓ decreases rapidly with the level ℓ , the cost C_ℓ for evaluating the differences $g_\ell = f_\ell - f_{\ell+1}$ is well approximated by \hat{C}_ℓ . The ratio of the total cost encountered when reducing the variance to a given value between multilevel Monte-Carlo (with optimal choice of N_ℓ) and standard Monte-Carlo (2.1) is then given by

$$\left(\sum_{\ell=1}^L \sqrt{V_\ell C_\ell} \right)^2 / (\mathbb{V}[f] C_1).$$

This is the basic quantitative relation indicating how the costs C_ℓ to evaluate the f_ℓ and the variances V_ℓ of the differences $f_\ell - f_{\ell+1}$ have to relate in order for the multilevel approach to be more efficient than standard Monte-Carlo estimation of f .

3. Stochastic estimation of the trace of a matrix. We now assume that we are given, in an indirect manner, a matrix $A = (a_{ij}) \in \mathbb{C}^{n \times n}$ for which we want to compute the trace

$$\mathrm{tr}(A) = \sum_{i=1}^n a_{ii}.$$

Our basic assumption is that the entries a_{ii} of A are neither available directly nor that they can all be obtained at decent computational cost. This is typically the case when A arises as a function of a large (and sparse) matrix, the most common case being the matrix inverse.

In a seminal paper [29], Hutchinson suggested to use a stochastic estimator to approximate $\mathrm{tr}(A)$. The following theorem summarizes his result together with the generalizations on the admissible probability spaces; see [13, 48], e.g.

THEOREM 3.1. *Let $\mathbb{P} : \Omega \rightarrow [0, 1]$ be a probability measure on a sample space Ω and assume that the components x_i of the vector $x \in \mathbb{C}^n$ are random variables depending on $\omega \in \Omega$ satisfying*

$$(3.1) \quad \mathbb{E}[x_i] = 0 \quad \text{and} \quad \mathbb{E}[\bar{x}_i x_j] = \delta_{ij} \quad (\text{where } \delta_{ij} \text{ is the Kronecker delta}).$$

Then

$$\mathbb{E}[x^* A x] = \mathrm{tr}(A) \quad \text{and} \quad \mathbb{V}[x^* A x] = \sum_{\substack{i,j,k,p=1 \\ i \neq j, k \neq p}}^n \bar{a}_{ij} a_{kp} \mathbb{E}[\bar{x}_i x_j \bar{x}_k x_p].$$

In particular, if the probability space is such that each component x_i is independent from x_j for $i \neq j$, then

$$\mathbb{V}[x^* A x] = \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ij} + \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ji} \mathbb{E}[x_i^2] \mathbb{E}[\bar{x}_j^2].$$

Proof. The proof is simple, but we repeat it here because the literature often treats only the real and not the general complex case. We have

$$\mathbb{E}[x^* A x] = \sum_{i=1}^n a_{ii} \mathbb{E}(\bar{x}_i x_i) + \sum_{i,j=1, i \neq j}^n a_{ij} \mathbb{E}(\bar{x}_i x_j) = \mathrm{tr}(A),$$

where the last inequality follows from (3.1). Similarly

$$\begin{aligned} \mathbb{V}[x^* A x] &= \mathbb{E} \left[\overline{(x^* A x - \mathrm{tr}(A))} (x^* A x - \mathrm{tr}(A)) \right] \\ &= \mathbb{E} \left[\left(\sum_{\substack{i,j=1 \\ i \neq j}}^n x_i \bar{a}_{ij} \bar{x}_j \right) \left(\sum_{\substack{k,p=1 \\ k \neq p}}^n \bar{x}_k a_{kp} x_p \right) \right] \\ (3.2) \quad &= \mathbb{E} \left[\sum_{\substack{i,j,k,p=1 \\ i \neq j, k \neq p}}^n \bar{a}_{ij} a_{kp} x_i \bar{x}_j \bar{x}_k x_p \right] = \sum_{\substack{i,j,k,p=1 \\ i \neq j, k \neq p}}^n \bar{a}_{ij} a_{kp} \mathbb{E}[x_i \bar{x}_j \bar{x}_k x_p]. \end{aligned}$$

Since the components x_i are assumed to be independent, we have $\mathbb{E}[\bar{x}_i x_j \bar{x}_k x_p] = 0$ except when $i = j, k = p$ (which does not occur in (3.2)) or $i = k, j = p$ or $i = p, j = k$.

This gives

$$\sum_{\substack{i,j,k,p=1 \\ i \neq j, k \neq p}}^n \bar{a}_{ij} a_{kp} \mathbb{E}[x_i \bar{x}_j \bar{x}_k x_p] = \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ij} \mathbb{E}[x_i \bar{x}_j \bar{x}_i x_j] + \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ji} \mathbb{E}[x_i \bar{x}_j \bar{x}_j x_i],$$

and in the first sum $\mathbb{E}[x_i \bar{x}_j \bar{x}_i x_i] = \mathbb{E}[\bar{x}_i x_i] \mathbb{E}[\bar{x}_j x_j] = 1$ by assumption, whereas in the second sum we have $\mathbb{E}[x_i \bar{x}_j \bar{x}_j x_i] = \mathbb{E}[x_i^2] \mathbb{E}[\bar{x}_j^2]$. \square

Note that as a definition for the variance of a complex random variable y we used $\mathbb{E}[(y - E(y))(y - E[y])]^2$ rather than $\mathbb{E}[(y - E[y])^2]$ to keep it real and non-negative. Standard choices for the probability spaces are to take x with identically and independently distributed (i.i.d.) components as

$$(3.3) \quad x_i \in \{-1, 1\} \text{ with equal probability } \frac{1}{2},$$

$$(3.4) \quad x_i \in \{-1, 1, -i, i\} \text{ with equal probability } \frac{1}{4},$$

$$(3.5) \quad x_i = \exp(i\theta) \text{ with } \theta \text{ uniformly distributed in } [0, 2\pi],$$

$$(3.6) \quad x_i \text{ is } N(0, 1) \text{ normally distributed.}$$

COROLLARY 3.2. *If the components x_i are i.i.d. with distribution (3.3) or (3.6), then*

$$\mathbb{V}[x^* A X x] = \frac{1}{2} \|\text{offdiag}(A + A^T)\|_F^2,$$

where $\|\cdot\|_F$ denotes the Frobenius norm and offdiag the offdiagonal part of a matrix. If the components are i.i.d. with distribution (3.4) or (3.5), then

$$\mathbb{V}[x^* A X x] = \|\text{offdiag}(A)\|_F^2.$$

Proof. For the distributions (3.3) and (3.6), the components x_i have only real values and $\mathbb{E}[x_i^2] = 1$. Therefore

$$\begin{aligned} \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ij} + \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ji} \mathbb{E}[x_i^2] \mathbb{E}[\bar{x}_j^2] &= \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ij} + \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ji} \\ &= \frac{1}{2} \sum_{\substack{i,j \\ i \neq j}}^n (\bar{a}_{ij} + a_{ji})(a_{ij} + a_{ji}) \\ &= \frac{1}{2} \|\text{offdiag}(A + A^T)\|_F^2. \end{aligned}$$

For the distributions (3.4) and (3.5) we have $\mathbb{E}[x_i^2] = 0$, and thus

$$\sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ij} + \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ji} \mathbb{E}[x_i^2] \mathbb{E}[\bar{x}_j^2] = \sum_{\substack{i,j \\ i \neq j}}^n \bar{a}_{ij} a_{ij} = \|\text{offdiag}(A)\|_F^2. \quad \square$$

In a practical situation where we approximate $\text{tr}(A)$ by averaging over N samples we can compute the root mean square deviation along with the averages and rely on the law of large numbers to assess the probability that the computed mean lies within the σ , 2σ or 3σ interval. Several results on Hutchinson's method have been formulated which go beyond this asymptotic aspects by giving tail or concentration bounds; see [2, 11, 37], e.g. For the sake of illustration we here report a summary of these results as given in [33]. In our numerical examples, we will simply work with the root mean square deviation to assess accuracy.

THEOREM 3.3. *Let the distribution for the i.i.d. components of the random vectors x^i be sub-Gaussian, and let $\epsilon, \delta \in (0, 1)$. Then for $N = \mathcal{O}(\log(1/\delta)/\epsilon^2)$ we have that the probability for*

$$(3.7) \quad \left| \frac{1}{N} \sum_{i=1}^n (x^i)^* A x^i - \text{tr}(A) \right| \leq \epsilon \|A\|_F$$

is $\geq 1 - \delta$.

Note that if A is symmetric positive semidefinite with λ_i denoting its (non-negative) eigenvalues, then

$$\|A\|_F = \left(\sum_{i=1}^n \lambda_i^2 \right)^{1/2} \leq \sum_{i=1}^n \lambda_i = \text{tr}(A),$$

implying that (3.7) yields a (probabilistic) relative error bound for the trace. Also note that the real distributions (3.3) and (3.6) are sub-Gaussian; see [33].

4. Multilevel Monte-Carlo for the trace of the inverse. We now turn to the situation where we want to estimate $\text{tr}(A^{-1})$ for a large and sparse matrix A . Direct application of Theorem 3.1 shows that an unbiased estimator for $\text{tr}(A^{-1})$ is given by

$$(4.1) \quad \frac{1}{N} \sum_{i=1}^N x^{(i)} A^{-1} x^{(i)} \approx \text{tr}(A^{-1}),$$

where the vectors $x^{(i)}$ are independent random variables satisfying (3.1), and that its variance is

$$\frac{1}{N} \|\text{offdiag}(A^{-1} + A^{-T})\|_F^2 \quad \text{or} \quad \frac{1}{N} \|\text{offdiag}(A^{-1})\|_F^2,$$

depending on whether the components of $x^{(i)}$ satisfy (3.3), (3.6) or (3.4), (3.5), respectively.

Each time we add a sample i to (4.1) we have to solve a linear system with matrix A and right hand side $x^{(i)}$, and the cost for solving these linear systems determines the cost for each stochastic estimate. For a large class of matrices, multigrid methods represent particularly efficient linear solvers. We assume that this is the case for our matrix A and now describe how to derive a multilevel Monte-Carlo method for the approximation of $\text{tr}(A^{-1})$ which uses the multigrid hierarchy not only for the linear solver, but also to obtain a good representation (2.3) required for a multilevel Monte-Carlo approach.

4.1. Derivation of a multilevel Monte-Carlo method. Multigrid methods rely on the interplay between a smoothing iteration and a coarse grid correction which are applied alternately. In the geometric interpretation, where we view components of vectors as representing a continuous function on a discrete grid, the smoother has the property to make the error of the current iterate smooth, i.e. varying slowly from one grid point to the next. Such error can be represented accurately by a coarser grid, and the coarse grid correction solves for this coarse error on the coarse grid using a coarse grid representation of the matrix. The solution is then interpolated back to the original “fine” grid and applied as a correction to the iterate. The principle can be applied recursively using a sequence of coarser grids with corresponding operators, the solves on the coarsest grid being obtained by direct factorization.

To obtain a multilevel Monte-Carlo decomposition we discard the smoother and only consider the coarse grid operators and the intergrid transfer operators which we now describe algebraically. The coarse grid operators are given by a sequence of matrices

$$A_\ell \in \mathbb{C}^{n_\ell \times n_\ell}, \ell = 1, \dots, L,$$

representing the original matrix $A = A_1 \in \mathbb{C}^{n_1 \times n_1}$ on the different levels $\ell = 1, \dots, L$; the prolongation and restriction operators

$$P_\ell \in \mathbb{C}^{n_\ell \times n_{\ell+1}}, R_\ell \in \mathbb{C}^{n_{\ell+1} \times n_\ell}, \ell = 1, \dots, L-1.$$

transfer data between the levels. Typically, when A is Hermitian, one takes $P_\ell = R_\ell^*$, and for given P_ℓ, R_ℓ the coarse system matrices A_ℓ are often constructed using the Petrov-Galerkin approach

$$A_{\ell+1} = R_\ell A_\ell P_\ell, \ell = 1, \dots, L-1.$$

Using the accumulated prolongation and restriction operators

$$\hat{P}_\ell = P_1 \cdots P_{\ell-1} \in \mathbb{C}^{n \times n_\ell}, \hat{R}_\ell = R_{\ell-1} \cdots R_1 \in \mathbb{C}^{n_\ell \times n}, \ell = 1, \dots, L,$$

where we put $\hat{R}_1 = \hat{P}_1 = I \in \mathbb{C}^{n \times n}$ by convention, we regard $\hat{P}_\ell A_\ell^{-1} \hat{R}_\ell$ as the approximation to A^{-1} at level ℓ . We thus obtain a multilevel decomposition for the trace as

$$(4.2) \quad \text{tr}(A^{-1}) = \sum_{\ell=1}^{L-1} \text{tr} \left(\hat{P}_\ell A_\ell^{-1} \hat{R}_\ell - \hat{P}_{\ell+1} A_{\ell+1}^{-1} \hat{R}_{\ell+1} \right) + \text{tr}(\hat{P}_L A_L^{-1} \hat{R}_L).$$

This gives

$$\text{tr}(A^{-1}) = \sum_{\ell=1}^{L-1} \mathbb{E} \left[(x^\ell)^* \left(\hat{P}_\ell A_\ell^{-1} \hat{R}_\ell - \hat{P}_{\ell+1} A_{\ell+1}^{-1} \hat{R}_{\ell+1} \right) x^\ell \right] + \mathbb{E} \left[(x^L)^* \hat{P}_L A_L^{-1} \hat{R}_L x^L \right],$$

with the components of $x^\ell \in \mathbb{C}^n$ being i.i.d. stochastic variables satisfying (3.1). The unbiased multilevel Monte-Carlo estimator is then

$$\begin{aligned} \text{tr}(A^{-1}) &\approx \sum_{\ell=1}^{L-1} \sum_{i=1}^{N_\ell} \left((x^{i,\ell})^* \hat{P}_\ell A_\ell^{-1} \hat{R}_\ell x^{i,\ell} - (x^{i,\ell})^* \hat{P}_{\ell+1} A_{\ell+1}^{-1} \hat{R}_{\ell+1} x^{i,\ell} \right) \\ &\quad + \sum_{i=1}^{N_L} (x^{i,L})^* \hat{P}_L A_L^{-1} \hat{R}_L x^{i,L}, \end{aligned}$$

where the vectors $x^{i,\ell} \in \mathbb{C}^n$ are stochastically independent samples of the random variable $x \in \mathbb{C}^n$ satisfying (3.1).

The following remarks collect some important observations about this stochastic estimator.

Remark 4.1. Computationally, the estimator requires to solve systems of the form $A_\ell y^{n,\ell} = z$ with $z = \hat{R}_\ell x^{n,\ell}$. Since the matrices A_ℓ arise from the multigrid hierarchy, we directly have a multigrid method available for these systems by restricting the method for A to the levels ℓ, \dots, L .

Remark 4.2. Since for any two matrices $B = (b_{ij}) \in \mathbb{C}^{n \times m}$ and $C = (c_{kl}) \in \mathbb{C}^{m \times n}$ the trace of their product does not depend on the order,

$$(4.3) \quad \operatorname{tr}(BC) = \sum_{i=1}^n \sum_{j=1}^m b_{ij} c_{ji} = \sum_{j=1}^m \sum_{i=1}^n c_{ji} b_{ij} = \operatorname{tr}(CB),$$

we have

$$\operatorname{tr}(\hat{P}_L A_L^{-1} \hat{R}_L) = \operatorname{tr}(A_L^{-1} \hat{P}_L \hat{R}_L).$$

So, instead of estimating the contribution $\operatorname{tr}(\hat{P}_L A_L^{-1} \hat{R}_L)$ in (4.2) stochastically, we can also compute it directly by inverting the matrix A_L and computing the product $A_L^{-1} \hat{R}_L \hat{P}_L$. Note that \hat{R}_L and \hat{P}_L are usually sparse with a maximum of d , say, non-zero entries per row. The arithmetic work for $A_L^{-1} \hat{R}_L \hat{P}_L$ is thus of order $\mathcal{O}(dN_L^2)$ for the product $\hat{R}_L \hat{P}_L$ plus $\mathcal{O}(n_L^3)$ for the inversion of A_L and the product $A_L^{-1}(\hat{R}_L \hat{P}_L)$. Since the variance of $x^* \hat{P}_L A_L^{-1} \hat{R}_L x$ is presumably large, this direct computation can be much more efficient than a stochastic estimation, even when we aim at only quite low precision in the stochastic estimate.

Remark 4.3. There are situations where $\hat{R}_\ell \hat{P}_\ell = I \in \mathbb{C}^{n_\ell \times n_\ell}$, for example in aggregation based multigrid methods, where the columns of P_ℓ are orthonormal and $R_\ell = P_\ell^*$, see [8, 9]. Then

$$\operatorname{tr}(\hat{P}_\ell A_\ell^{-1} \hat{R}_\ell) = \operatorname{tr}(A_\ell^{-1} \hat{R}_\ell \hat{P}_\ell) = \operatorname{tr}(A_\ell^{-1}),$$

and

$$\operatorname{tr}(\hat{P}_{\ell+1} A_{\ell+1}^{-1} \hat{R}_{\ell+1}) = \operatorname{tr}(\hat{P}_\ell P_\ell A_{\ell+1}^{-1} R_\ell \hat{R}_\ell) = \operatorname{tr}(P_\ell A_{\ell+1}^{-1} R_\ell \hat{R}_\ell \hat{P}_\ell) = \operatorname{tr}(P_\ell A_{\ell+1}^{-1} R_\ell).$$

This means that instead of the multilevel decomposition (4.2) we can use

$$\operatorname{tr}(A) = \sum_{\ell=1}^{L-1} (\operatorname{tr}(A_\ell^{-1}) - \operatorname{tr}(P_\ell A_{\ell+1}^{-1} R_\ell)) + \operatorname{tr}(A_L^{-1}),$$

in which the stochastic estimation on level ℓ now involves random vectors from \mathbb{C}^{n_ℓ} instead of \mathbb{C}^n .

4.2. Discussion of the multilevel Monte-Carlo method. A profound analysis of the proposed multilevel Monte-Carlo method must take the approximation properties of the representation of the matrix at the various levels into account. This is highly problem dependent, so that in this paper we only provide a discussion of heuristics on why the proposed approach has the potential to yield efficient multilevel Monte-Carlo schemes.

To simplify the discussion to follow, let us assume that the variance of the estimator at level ℓ is given by the square of the Frobenius norm of the off-diagonal part. This is the case, for example, if the components are i.i.d. with distribution (3.4) or (3.5); see Corollary 3.2. This Frobenius norm can be related to the singular values of A . Recall that the singular value decomposition of a non-singular matrix A is

$$(4.4) \quad \begin{aligned} A &= U \Sigma V^* \quad \text{with } U, \Sigma, V \in \mathbb{C}^{n \times n}, U^* U = V^* V = I, \\ U &= [u_1 | \dots | u_n], \quad V = [v_1 | \dots | v_n], \\ \Sigma &= \operatorname{diag}(\sigma_1, \dots, \sigma_n), \quad 0 < \sigma_1 \leq \dots \leq \sigma_n, \end{aligned}$$

with left singular vectors u_i , right singular vectors v_i and positive singular values σ_i which we ordered by increasing value for convenience here. In the following we base all our discussion on singular values and vectors. It is therefore worthwhile to mention that in the case of a Hermitian matrix A this discussion simplifies in the sense that then the singular values are the moduli of the eigenvalues, and left and right singular vectors are identical and coincide with the eigenvectors.

LEMMA 4.4. *Let $A \in \mathbb{C}^{n \times n}$ have singular values $\sigma_i, i = 1, \dots, n$. Then*

$$(4.5) \quad \|\text{offdiag}(A)\|_F^2 = \sum_{i=1}^n \sigma_i^2 - \sum_{i=1}^n |a_{ii}|^2.$$

Proof. The equality $\|A\|_F^2 = \sum_{i=1}^n \sigma_i^2$ is a basic fact from linear algebra, see [24], e.g. The formula (4.5) uses this and corrects for the vanishing diagonal part in $\text{offdiag}(A)$. \square

For the trace of the inverse A^{-1} we have

$$(4.6) \quad \|\text{offdiag}(A^{-1})\|_F^2 = \sum_{i=1}^n \sigma_i^{-2} - \sum_{i=1}^n |(A^{-1})_{ii}|^2.$$

since the reciprocals of the singular values of A , are the singular values of A^{-1} . Therefore, in a simplified manner—disregarding the second term in (4.6)—it appears that the small singular values of A are those who contribute most to the variance for the Hutchinson estimator (4.1) for $\text{tr}(A^{-1})$. In high performance computing practice, *deflation* has thus become a common tool, see [12, 14, 19, 22], e.g., to reduce the variance: One precomputes the k , say, smallest singular values $\sigma_1, \dots, \sigma_k$ of A in the singular value decomposition (4.4) together with their left singular vectors u_1, \dots, u_k . With the orthogonal projector

$$(4.7) \quad \Pi = U_k U_k^*, \text{ where } U_k = [u_1 | \dots | u_k],$$

we now have $A^{-1} = A^{-1}(I - \Pi) + A^{-1}\Pi$ with

$$(4.8) \quad A^{-1}(I - \Pi) = \sum_{i=k+1}^n v_i \sigma_i^{-1} u_i^*, \quad A^{-1}\Pi = \sum_{i=1}^k A^{-1} u_i u_i^*.$$

This shows that in $A^{-1}(I - \Pi)$ we have deflated the small singular values of A , so that we can expect a reduction of the variance when estimating the trace of this part stochastically. The trace of the second part is equal to the sum $\sum_{i=1}^k u_i^* A^{-1} u_i$ (see (4.3)), and $A^{-1} u_i = \sigma_i^{-1} v_i$. So the second part can be computed directly from the singular triplets computed for the deflation. If A is Hermitian, the deflation approach simplifies and amounts to precomputing the k smallest eigenpairs. We refer to the results in [19] for a more in-depth analysis and discussion about the heuristics just presented.

The deflation approach is still quite costly, since one has to precompute the singular values and vectors, and if the size of the matrix increases it is likely that we have to increase k to maintain the same reduction in the variance. Approximate deflation has thus been put forward as an alternative, see [4, 36], where one can use larger values for k while at the same time allowing that the contribution of the small singular values to the variance is eliminated only approximately. One then replaces Π by a more general projector of the form

$$\Pi = \hat{U}_k (\hat{V}_k^* A \hat{U}_k)^{-1} \hat{V}_k^* A, \quad \hat{U}_k, \hat{V}_k \in \mathbb{C}^{n \times k}$$

where now \hat{U}_k and \hat{V}_k can be regarded as containing approximate left and right singular vectors, respectively, as their columns. Actually, it is sufficient that their range is spanned by such approximations to left and right singular vectors, since the construction of Π is invariant under transformations $\hat{U} \rightarrow \hat{U}B_U, \hat{V} \rightarrow \hat{V}B_V$ with non-singular matrices $B_U, B_V \in \mathbb{C}^{k \times k}$. In the decomposition $A^{-1} = A^{-1}(I - \Pi) + A^{-1}\Pi$ we now have, again using (4.3),

$$\begin{aligned} \text{tr}(A^{-1}(I - \Pi)) &= \text{tr}(A^{-1}) - \text{tr}(\hat{U}_k(\hat{V}_k^* A \hat{U}_k)^{-1} \hat{V}_k^*), \\ \text{tr}(A^{-1}\Pi) &= \text{tr}(\hat{U}_k(\hat{V}_k^* A \hat{U}_k)^{-1} \hat{V}_k^*). \end{aligned}$$

If k is relatively small, the second trace can be computed directly as in the exact deflation approach. If we take larger values for k , we can estimate it stochastically. The inexact deflation approach then becomes a two-level Monte-Carlo method.

If we look at our multilevel Monte-Carlo decomposition (2.2) with just two levels, then it differs from inexact deflation in that the value for k is now very large, namely the grid size at level 2 which usually is $\mathcal{O}(n)$. The matrix \hat{U}_k spanning the approximate singular vectors is replaced by the prolongation operator P_1 , and \hat{V}_k^* corresponds to the restriction operator R_1 . The multigrid construction principle should ensure that the range of P_1 contains good approximations to $\mathcal{O}(n)$ left singular vectors belonging to small singular values, and similarly for R_1^* with respect to right singular vectors. This is why the variance reduction can be expected to be efficient. We thus have a large value of k —proportional to n —which targets at a high reduction of the variance of the first term. The second term involves the second level matrix representation, which is still of large size, and its trace estimator will, in addition, still have large variance. This is the reason why we extend the approach to involve many levels, ideally until a level L where we can compute the trace directly, so that we do not suffer from a potentially high variance of a stochastic estimator.

To conclude this discussion, we note that several other techniques for variance reduction have been suggested which can also be regarded as two-level Monte-Carlo techniques. For example, [32, 5] take a decomposition $A^{-1} - p(A) + p(A)$ with an appropriately chosen polynomial $p(A)$ and then estimates $\text{tr}(A^{-1} - p(A))$ stochastically. The “truncated solver” method of [1] follows a related idea by subtracting an approximation to the inverse. A similar decomposition with p being a truncated Chebyshev series approximation was considered in [26, 27, 47], for example, in which case $\text{tr}(A^{-1} - p(A))$ is actually neglected. The work then resides in the stochastic estimation of $\text{tr}(p(A))$, thus avoiding to solve linear systems.

Finally, we refer to [33] for a recent further variance reduction technique for Hutchinson’s method, enhancing it by using vectors of the form Av with random vectors v .

5. Numerical results. We consider three classes of matrices: The standard discrete 2d Laplacian, the 2d gauge Laplacian, and the Schwinger model. These three classes represent an increasingly complex hierarchy of problems which will eventually lead to our final, though yet unreached target, the Wilson-Dirac matrix arising in lattice QCD. The improvements of the multilevel approach compared to “plain” Hutchinson (4.1) are tremendous and typically reach two orders of magnitude or more. This is why we compare against *deflated Hutchinson*, where we deflate the n_{def} smallest eigenpairs of the matrix A . With $U \in \mathbb{C}^{n \times n_{\text{def}}}$ holding the respective eigenvectors in its columns, we use the projector $\Pi = I - UU^*$ as in (4.7), resulting in the decomposition (4.8) Therein we estimate $\text{tr}(A^{-1}(I - \Pi))$ with the Hutchinson estimator whereas $\text{tr}(A^{-1}\Pi) = \sum_{i=1}^{n_{\text{def}}} \lambda_i^{-1}$ is obtained directly from the deflated eigenpairs. We

always performed a rough scan to determine a number n_{def} of deflated eigenpairs which s close to time-optimal. The deflated Hutchinson approach usually gains at least one order of magnitude in time and arithmetic cost over plain Hutchinson.

All our computations were done on a single thread of an Intel Xeon Processor E5-2699 v4, with a MATLAB R2021a implementation of our numerical experiments for the 2d Laplacian, and in Python for our tests with the gauge Laplacian and the Schwinger model. By default we aimed at a relative accuracy of $\epsilon = 10^{-3}$. This is done as follows: We first perform five stochastic estimates, take their mean and subtract their root mean square deviation, giving the value τ . In the deflated Hutchinson method we now perform stochastic estimates as long as their root mean square deviation exceeds $\epsilon\tau$. For the multilevel Monte-Carlo method we have to prescribe a value for the root mean square deviations ρ_ℓ for the stochastic estimation of each of the traces

$$(5.1) \quad \text{tr} \left(\hat{P}_\ell A_\ell^{-1} \hat{R}_\ell - \hat{P}_{\ell+1} A_{\ell+1}^{-1} \hat{R}_{\ell+1} \right), \quad \ell = 1, \dots, L-1$$

from (4.2), while we always compute the last term $\text{tr}(\hat{P}_L A_L^{-1} \hat{R}_L)$ in (4.2) non-stochastically as $\text{tr}(A_L^{-1} \hat{R}_L \hat{P}_L)$, inverting A_L explicitly. The requirement is to have

$$\sum_{\ell=1}^{L-1} \rho_\ell^2 = (\epsilon\tau)^2,$$

so the obvious approach is to put $\rho_\ell = \epsilon\tau/\sqrt{L-1}$ for all ℓ , and this is what we do in our first two examples. It might be advantageous, though, to allow for a larger value of ρ_ℓ on those level differences where the cost is high, and we do so in Example 5.3. For each stochastic estimate for (5.1) we have to solve linear systems with the matrices A_ℓ and $A_{\ell+1}$. This is done using a multigrid method based on the same prolongations P_ℓ , restrictions R_ℓ and coarse grid operators A_ℓ that we use to obtain our multilevel decomposition (4.2). However, when multigrid is used as a solver, we use the full hierarchy going down to coarse grids of very small sizes, whereas in the multilevel decomposition (4.2) used in multilevel Monte-Carlo we might stop at an earlier level. For all experiments we report mainly two quantities. The first is the number of stochastic estimates that are performed at each level difference (5.1) for multilevel Monte-Carlo together with the number of stochastic estimates in deflated Hutchinson (which always require linear solves at the finest level). These numbers may be interpreted as illustrating how multilevel Monte-Carlo moves the higher variances to the coarser level differences. As a second quantity, we report the approximate arithmetic cost for both methods, deflated Hutchinson and multilevel Monte-Carlo, which we obtain using the following cost model: For every matrix-vector product of the form Bx we assume a cost of $\text{nnz}(B)$, the number of nonzeros in B . In this manner, one unit in the cost model roughly corresponds to a multiplication plus an addition. This applies to the computation of residuals, of prolongations and restrictions and the coarsest grid solve in the multigrid solver as well as to the “global” restrictions and prolongations $\hat{R}_\ell, \hat{P}_\ell$ used in each stochastic estimate in multilevel Monte-Carlo. For the latter method, we also count the work for the direct computation of the trace at the coarsest level, which involves the inversion of the coarsest grid matrix and additional matrix-matrix products. This cost model thus only neglects vector-vector and scalar operations and is thus considered sufficiently accurate for our purposes.

		2d Laplace							
N		$\ell = 1$	$\ell = 2$	$\ell = 3$	$\ell = 4$	$\ell = 5$	$\ell = 6$	n_{defl}	L
63	n_ℓ	63^2	31^2	15^2				92	3
	$\text{nnz}(L_\ell^N)$	1.96e4	8.28e3	1.85e3					
127	n_ℓ	127^2	63^2	31^2	15^2			44	4
	$\text{nnz}(L_\ell^N)$	8.01e4	3.50e4	8.28e3	1.85e3				
255	n_ℓ	255^2	127^2	63^2	31^2	15^2		64	5
	$\text{nnz}(L_\ell^N)$	3.24e5	1.44e5	3.50e4	8.28e3	1.85e3			
511	n_ℓ	511^2	255^2	127^2	63^2	31^2	15^2	76	6
	$\text{nnz}(L_\ell^N)$	1.30e6	5.82e5	1.44e5	3.50e4	8.28e3	1.85e3		

TABLE 1
Parameters and quantities for Example 5.1

Example 5.1. The discrete 2d Laplacian is the $N^2 \times N^2$ -matrix

$$L^N = B \otimes I + I \otimes B, \quad I \text{ the } N \times N \text{ identity, } B_0 = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & & -1 & 2 \end{bmatrix} \in \mathbb{C}^{N \times N},$$

which results from the finite difference approximation of the Laplace operator on an equidistant grid in the unit square with Dirichlet boundary conditions. Note that the eigenvalues of L^N are explicitly known, so the trace of the inverse is, in principle, directly available as the sum of the inverses of the eigenvalues.

For the multigrid hierarchy we choose P_ℓ to be the standard bilinear interpolation from a grid of size $N_{\ell+1} \times N_{\ell+1}$ to one of size $N_\ell \times N_\ell$; see [45]. Here, the number N_ℓ of grid points in one dimension on level ℓ is recursively given as $N_{\ell+1} = \lfloor N_\ell/2 \rfloor$. The restrictions P_ℓ are taken as the adjoints of the interpolations P_ℓ , and the coarse grid operators L_ℓ^N are obtained as Galerkin approximations $L_{\ell+1}^N = R_\ell L_\ell^N P_\ell$. So the operator L_ℓ^N at level ℓ is an $n_\ell \times n_\ell$ matrix with $n_\ell = N_\ell^2$.

V-cycle multigrid with one step of Gauss-Seidel pre- and post-smoothing was used as a solver to compute $(L_\ell^N)^{-1}x$ on the various levels. In the solver, we always use the full hierarchy down to a smallest coarsest grid operator size of $N^2 = 7^2$, where we inverted directly using a Cholesky factorization. For multilevel Monte-Carlo we took L , the maximum number of levels, such that $N_L = 15$, since it turned out that with this choice the work for the direct computation of the trace at level L —requiring the inversion of a matrix of size $15^2 = 225$ —was small enough compared to the other cost. Table 1 summarizes the most important quantities for the matrices in this example. Figure 1 reports timings and arithmetic cost for a study to approximately find the time-optimal number of eigenvalues to deflate in the deflated Hutchinson method for $N = 127$. The left part of the figure shows that there is a substantial decrease in the timings when we start to increase the number of deflated eigenvalues up to around 50, after which the cost for computing the eigenpairs becomes increasingly larger, eventually dominating the overall execution time completely. The straight horizontal line on the bottom indicates the time required in the multilevel Monte-Carlo approach using $L = 4$ levels. Even with an optimal number of deflated eigenvalues, deflated Hutchinson takes about 7 times longer than multilevel Monte-Carlo. The right part of Figure 1 reports the arithmetic cost. We do *not* include the cost for the computation of the eigenpairs, since we ignore the details of Matlab's `eigs` function that we use

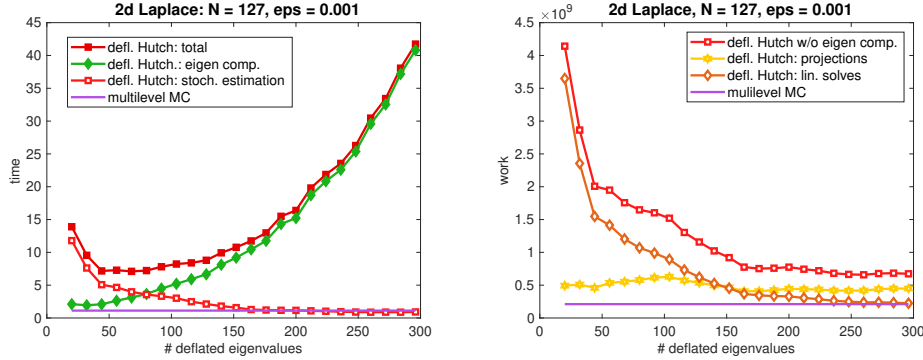


FIG. 1. *2d Laplace: Comparison of execution times and arithmetic cost for multilevel Monte-Carlo and deflated Hutchinson, varying n_{defl} .*

to compute those. The top line in the right part of the figure thus corresponds to the line marked with open squares in the left part, and we see that the cost model quite accurately matches the observed execution times. The two lines below the top line separate this work into the work spent in the computation of the projections $(I - UU^*)x = x - U(U^*x)$, which is $2nn_{\text{defl}}$, and the work spent in the linear solves. The straight horizontal line represents the work for multilevel Monte-Carlo. We see that if a large number of eigenpairs is deflated, we have comparable work in the linear solves in deflated Hutchinson than in multilevel Monte-Carlo, due to the fact that we have to do a similar number of stochastic estimates. Each estimate, though, now becomes substantially more expensive due to the additional projections, so that we still see a factor of about 3 in favor of multilevel Monte-Carlo. We repeat that this comparison does not take the work for computing the eigenpairs in deflated Hutchinson into account.

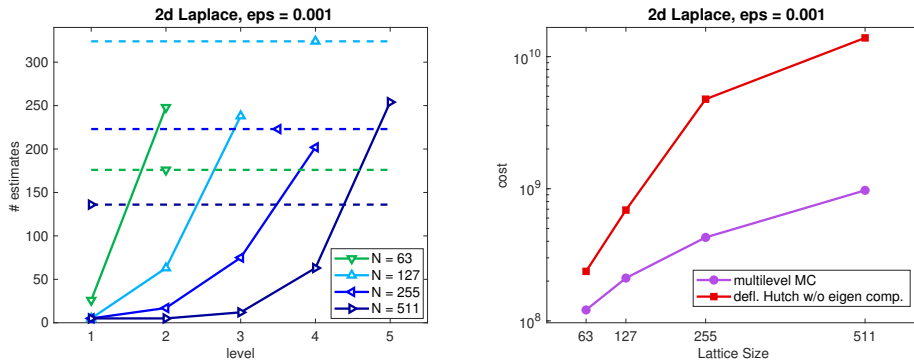


FIG. 2. *Comparison of multilevel Monte-Carlo and optimally deflated Hutchinson for the 2d Laplace matrix: no of stochastic estimates on each level difference (5.1) and total work for different N .*

We proceed with Figure 2 which reports results for the four size parameters $N = 63, 127, 255, 511$. For multilevel Monte-Carlo, the left part of the figure shows the number of stochastic estimates that we perform for each of the differences $\hat{P}_\ell(L_\ell^N)^{-1}\hat{R}_\ell - \hat{P}_{\ell+1}(L_{\ell+1}^N)^{-1}\hat{R}_{\ell+1}$ in (5.1). For comparison, the number of stochastic estimates required in (time optimally) deflated Hutchinson are indicated as dashed horizontal

lines; see Table 1 for the values of n_{defl} that we used. The right part of Figure 2 shows the total amount of work invested in the different methods where, again, we do not count the work for the eigenpair computation in deflated Hutchinson. The figure illustrates the fact that in multilevel Monte-Carlo we do few estimates on the fine, expensive levels and more on the coarse and cheap levels. The plots for the cost in the right part show how this translates into a reduction of the arithmetic cost, reaching about one order of magnitude for the larger values of N , even without accounting for the cost for the computation of the eigenpairs in deflated Hutchinson.

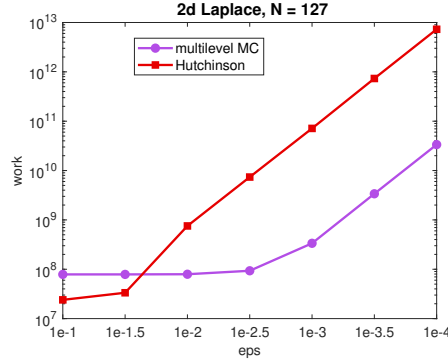


FIG. 3. Arithmetic cost for multilevel Monte-Carlo and standard Hutchinson as a function of the target accuracy ϵ .

Our multilevel Monte-Carlo approach performs a Monte-Carlo estimate for each level difference. It therefore exhibits the same quadratic scaling with respect to the target accuracy ϵ as standard Hutchinson does. This is illustrated in Figure 3 for $N = 127$. This is a logscale plot, and the quadratic dependence on ϵ^{-1} is clearly visible for the smaller values of ϵ . For the larger values of ϵ (to the left), the results are less significant, since we always do a minimum of at least 5 stochastic estimates for each level difference. For $\epsilon = 10^{-1}$ and $\epsilon = 10^{-1.5}$, this is precisely what happens for all the level differences, so the cost for multilevel Monte-Carlo is exactly the same (and probably unnecessarily high) for both these values of ϵ . Note that the arithmetic cost of standard Hutchinson is between 2 and 3 orders of magnitude higher than that of multilevel Monte-Carlo for $\epsilon \leq 10^{-2}$.

Example 5.2. The gauge Laplacian G^N is a modification of the standard 2d discrete Laplace matrix where now the coupling coefficients—called gauge links—are complex numbers of modulus one but with a random phase. The gauge Laplacian represents a first step towards the modeling of gauge field theories in physics, where the random coefficients model the fluctuating background gauge field. With u_{ij} describing the variables belonging to a grid point (ih, jh) , $h = 1/N$, $i, j = 0, \dots, N - 1$, the coupling described by one row of the gauge Laplacian $G^N \in \mathbb{C}^{N^2 \times N^2}$ is given as

$$4u_{ij} - e^{i\Theta_{ij}} u_{i+1,j} - e^{i\Phi_{ij}} u_{i,j+1} - e^{-i\Theta_{i-1,j}} u_{i-1,j} - e^{-i\Phi_{i,j-1}} u_{i,j-1},$$

$$i, j = 0, \dots, N - 1,$$

where the indices are to be understood mod N since we have periodic boundary conditions. Gauge Laplacians are Hermitian and positive semidefinite. Typically, as in our case, they are even positive definite.

		Gauge Laplace					
N		$\ell = 1$	$\ell = 2$	$\ell = 3$	$\ell = 4$	n_{def}	L
64	n_ℓ	4096	1354	134		60	3
	$\text{nnz}(G_\ell^N)$	20480	24900	3172			
128	n_ℓ	16384	5440	554		60	3
	$\text{nnz}(G_\ell^N)$	81920	99448	11300			
256	n_ℓ	65536	21802	2348	196	20	4
	$\text{nnz}(G_\ell^N)$	327680	394628	49416	6352		
512	n_ℓ	262144	87296	9562	924	20	4
	$\text{nnz}(G_\ell^N)$	327680	394628	49416	6352		

TABLE 2
Parameters and quantities for Example 5.2

The Python package pyAMG, see [34], contains functions to produce gauge Laplacians with prescribed distributions for the phases Θ_{ij} and Φ_{ij} of the gauge links. For our experiments, we produced gauge Laplacians with grid spacing $a = 1$ and temperature $\beta = 0.009$ for the distribution of the phases of the gauge links.

The pyAMG package provides a variety of algebraic multigrid methods. To build a multigrid hierarchy for gauge Laplacians, we use adaptive smoothed aggregation, where for the adaptive setup we took the parameters

$$\text{num_candidates} = 2, \text{ candidate_iters} = 5, \text{ improvement_iters} = 8.$$

As our linear solver we used V-cycle multigrid with one step of Gauss-Seidel pre- and post-smoothing. Table 2 summarizes the most important quantities for the same four lattice sizes as we had for the 2d Laplacian. The table shows that smoothed aggregation yields matrices at level 2 which are a factor 3 less in size but are significantly more dense, since they have even more non-zeros than the matrices at level 1. For the subsequent levels the coarsening is quite aggressive, reducing the sizes by factors of roughly 10, and similarly for the number of non-zeros.

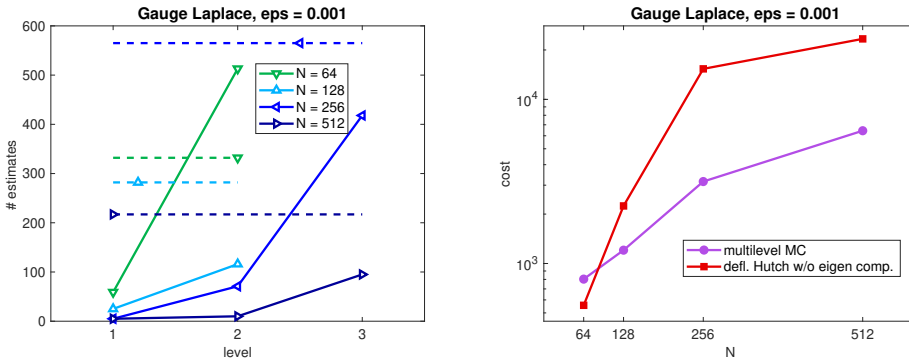


FIG. 4. Comparison of multilevel Monte-Carlo and optimally deflated Hutchinson for the gauge Laplace matrices: no of stochastic estimates on each level difference (5.1) and total work for different N .

In the same manner as Figure 2 did for the 2d Laplacian, Figure 4 now reports the number of estimates at the various levels and the arithmetic cost for multilevel Monte-Carlo as well as for optimally deflated Hutchinson. As for the first example the figure

illustrates that multilevel Monte-Carlo allows to do few estimates on the fine levels. The gain in work reaches up to a factor of 5 for the larger values of N , whereas $N = 64$ is too small in this example for multilevel Monte-Carlo to beat deflated Hutchinson. As before, we do not count the arithmetic cost for computing the eigenpairs in deflated Hutchinson in this comparison.

Example 5.3. Our third example is the Schwinger discretization of the 2-dimensional Dirac operator [41]. This operator describes quantum electrodynamics (QED), the quantum field theory of the electro-magnetic interaction between charged particles via photons. The Schwinger matrix resembles the gauge Laplacian in the sense that there is a periodic nearest neighbor coupling on an equidistant grid in the unit square and that the coupling coefficients are based on complex numbers of modulus 1 with a random phase. The difference is that for Dirac operators there are several (here: two) variables per grid point, representing different spins. With the Pauli matrices

$$\sigma_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \sigma_2 = \begin{bmatrix} 0 & i \\ -i & 0 \end{bmatrix},$$

and the understanding that a grid variable u_{ij} has now two components, i.e. $u_{ij} \in \mathbb{C}^2$, representing the different spins at grid point (ih, jh) , the periodic couplings is now given as

$$(5.2) \quad (4 + m) \cdot u_{ij} - e^{i\Theta_{ij}}(I - \sigma_1)u_{i+1,j} - e^{i\Phi_{ij}}(I - \sigma_2)u_{i,j+1}$$

$$(5.3) \quad - e^{-i\Theta_{i-1,j}}(I + \sigma_1)u_{i-1,j} - e^{-i\Phi_{i,j-1}}(I + \sigma_2)u_{i,j-1},$$

$$(5.4) \quad i, j = 1, \dots, N.$$

Note that the Pauli matrices cross-couple the spins. Thus, if we order spin components such that the first spin component at each grid location comes first and all second spin components follow, the Schwinger matrix has the form

$$(5.5) \quad S^N = \begin{bmatrix} G^N & B \\ -B^* & G^N \end{bmatrix}$$

where the matrices G^N are gauge Laplacians and the matrix B represents the spin cross-coupling.

We used a Schwinger matrix arising from a thermalized configuration within a Markov process. This guarantees that the random gauge links obey a Boltzmann distribution with a given temperature parameter. The matrix belongs to an $N \times N$ lattice with $N = 128$, and is thus of size $2N^2 \times 2N^2 = 32,768 \times 32,768$.

The multigrid hierarchy for the Schwinger matrix is obtained through an aggregation based approach similar to the one typically used for the 4-dimensional Wilson-Dirac operator; see [3, 17]. Giving all details would be beyond the scope of this paper, so here is a rough sketch: At each level, the operator represents a periodic nearest neighbor coupling on a 2-dimensional lattice of decreasing size. At each lattice site we have several, d say, degrees of freedom (dofs), i.e. variables belonging to a lattice site are vectors of length d . When going from one level to the next, we subdivide the lattice into small sublattices—the aggregates. Each aggregate becomes a single lattice site on the next level. The corresponding restriction operator is obtained by computing (quite inexact) approximations to the d smallest eigenvectors, the components of which are assembled over the aggregates and orthogonalized. This gives restriction operators which are orthonormal, and since we take the prolongations to be their adjoints, we

Schwinger model						
N		$\ell = 1$	$\ell = 2$	$\ell = 3$	$\ell = 4$	L
128	n_ℓ	$2 \cdot 128^2$	$4 \cdot 32^2$	$8 \cdot 8^2$	$8 \cdot 2^2$	4
	$\text{nnz}(S_\ell^N)$	2.94e5	1.64e5	2.46e4	1024	
	m	-0.1320	-0.1325	-0.1329	-0.1332	-0.1333
	n_{defl}	384	384	512	512	512

TABLE 3
Parameters and quantities for Example 5.3

are in the simplified situation of Remark 4.3 for estimating the traces of the differences in multilevel Monte-Carlo.

The Schwinger matrix is not Hermitian, but its eigenvalues come in complex conjugate pairs. This is due to a non-trivial symmetry induced by the spin structure that can be seen from (5.5),

$$JS^N = (S^N)^* J, \quad \text{where } J = \begin{bmatrix} I & 0 \\ 0 & -I \end{bmatrix},$$

so that to each right eigenpair (x, λ) of S^N corresponds a left eigenpair $((Jx)^*, \bar{\lambda})$. This spin symmetry can be preserved on the coarser levels if one doubles the dofs; see [3, 17].

We built a multigrid hierarchy with four levels. For the aggregates, at all levels we always aggregated 4×4 sublattices into one lattice site on the next level, and we started with 2 dofs for the second level and 4 for all remaining levels. Those dofs are then doubled because we implemented the spin structure preserving approach. Table 3 summarizes the most important information on the multigrid hierarchy. It also shows the five different (negative) values for the mass m that we used in our experiments. These values are physically meaningful, and for all of them the spectrum of S^N is contained in the right half plane. As m becomes smaller, S^N becomes more ill-conditioned, so the work for each stochastic estimate increases. When solving linear systems at the various levels, we used one V-cycle of multigrid with two steps of Gauss-Seidel pre- and post-smoothing as a preconditioner for flexible GMRES [39]. Our implementation was done in Python¹.

Figure 5 shows our results. We tuned the required root mean square deviation ρ_ℓ at each level due to the observation that this time the root mean square deviation is comparably small on the last level difference. The values we chose, independently of the mass parameter m , are $\rho_1 = \sqrt{0.4}\epsilon\tau$, $\rho_2 = \sqrt{0.55}\epsilon\tau$ and $\rho_3 = \sqrt{0.05}\epsilon\tau$ for all masses.

As in the other examples, we compare against deflated Hutchinson with a time-optimal number of deflated eigenpairs, and we do not count the work for the eigenpair computation. The figure shows that multilevel Monte-Carlo becomes increasingly more efficient over deflated Hutchinson as the masses become smaller, ending up in a one order of magnitude improvement in work for the smallest. Interestingly, we also see that the number of stochastic estimates to be performed on each level in multilevel Monte-Carlo depends on the masses only very mildly, whereas the number of stochastic estimates increase rapidly in deflated Hutchinson.

¹The programs can be found in the GitHub repository <https://github.com/Gustavroot/MLMCTraceComputer>

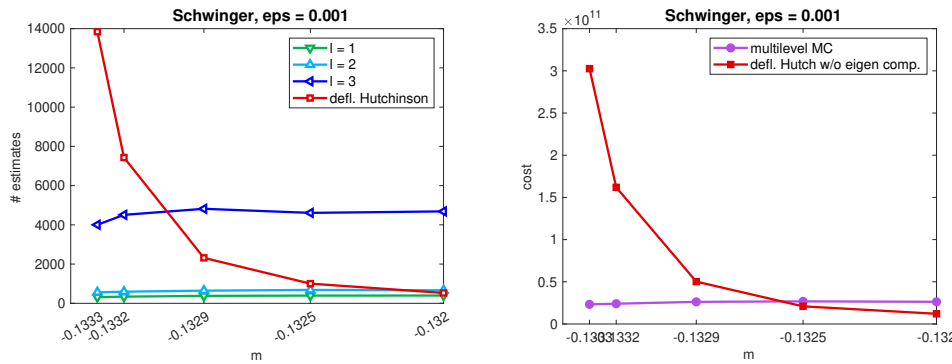


FIG. 5. Multilevel Monte-Carlo and deflated Hutchinson for the Schwinger matrix: no of stochastic estimates on each level difference (5.1) and total work for different masses m .

Conclusion. We presented a novel multilevel Monte-Carlo approach to stochastically estimate the trace of a matrix. The method relies on the availability of a multigrid hierarchy and estimates traces of differences of matrices at the various levels. The method is efficient if the variances at the earlier, finer level differences, are small, and this is what we observed in three different examples, two of which used an adaptive algebraic approach to determine the multigrid hierarchy.

Acknowledgments. We would like to thank Francesco Knechtli and Tomasz Korzec for stimulating discussion including references on work done in the lattice QCD community, and Karsten Kahl for providing us with the thermalized Schwinger model matrix.

REFERENCES

- [1] C. ALEXANDROU, M. CONSTANTINO, V. DRACH, K. HADJIYIANNAKOU, K. JANSEN, G. KOUTSO, A. STRELCHENKO, AND A. VAQUERO, *Evaluation of disconnected quark loops for hadron structure using GPUs*, Comput. Phys. Commun., 185 (2014), p. 13701382, <https://doi.org/10.1016/j.cpc.2014.01.009>, <http://dx.doi.org/10.1016/j.cpc.2014.01.009>.
- [2] H. AVRON AND S. TOLEDO, *Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix*, J. ACM, 58 (2011).
- [3] R. BABICH, J. BRANNICK, R. C. BROWER, M. A. CLARK, T. A. MANTEUFFEL, S. F. MCCORMICK, J. C. OSBORN, AND C. REBBI, *Adaptive multigrid algorithm for the lattice Wilson-Dirac operator*, Phys. Rev. Lett. 105:201602, (2010). arXiv:1005.3043v2 [hep-lat].
- [4] G. BALI, S. COLLINS, A. FROMMER, K. KAHL, I. KANAMORI, B. MÜLLER, M. ROTTMANN, AND J. SIMETH, *(approximate) low-mode averaging with a new multigrid eigensolver*, 2015, <https://arxiv.org/abs/1509.06865>.
- [5] S. BARAL, T. WHYTE, W. WILCOX, AND R. B. MORGAN, *Disconnected loop subtraction methods in lattice qcd*, Comput. Phys. Commun., 241 (2019), pp. 64–79.
- [6] C. BEKAS, E. KOKIOPOULOU, AND Y. SAAD, *An estimator for the diagonal of a matrix*, Applied Numerical Mathematics, 57 (2007), pp. 1214–1229.
- [7] A. H. BENTBIB, M. EL GHOMARI, K. JBILOU, AND L. REICHEL, *Shifted extended global Lanczos processes for trace estimation with application to network analysis*, Calcolo, 58 (2021), pp. Paper No. 4, 35, <https://doi.org/10.1007/s10092-020-00395-1>, <https://doi.org/10.1007/s10092-020-00395-1>.
- [8] D. BRAESS, *Towards algebraic multigrid for elliptic problems of second order*, Computing, 55 (1995), pp. 379–393, <https://doi.org/10.1007/BF02238488>.
- [9] M. BREZINA, R. FALGOUT, S. MACLACHLAN, T. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *Adaptive smoothed aggregation (α SA) multigrid*, SIAM Review, 47 (2005), pp. 317–346.
- [10] J. CHEN AND Y. SAAD, *A posteriori error estimate for computing $\text{tr}(f(A))$ by using the Lanczos method*, Numer. Linear Algebra Appl., 25 (2018), pp. e2170, 20, <https://doi.org/10.1002/>

- nla.2170, <https://doi.org/10.1002/nla.2170>.
- [11] A. CORTINOVIS AND D. KRESSNER, *On randomized trace estimates for indefinite matrices with an application to determinants*, 2020, <https://arxiv.org/abs/2005.10009>.
 - [12] T. A. DEGRAND AND S. SCHAEFER, *Improving Meson Two Point Functions in Lattice QCD*, *Comput. Phys. Commun.*, 159 (2004), pp. 185–191, <https://doi.org/10.1016/j.cpc.2004.02.006>, <https://arxiv.org/abs/hep-lat/0401011>.
 - [13] S. DONG AND K. LIU, *Stochastic estimation with Z_2 noise*, *Phys. Lett. B*, 328 (1994), pp. 130–136.
 - [14] E. ENDRESS, C. PENA, AND K. SIVALINGAM, *Variance Reduction with Practical All-to-all Lattice Propagators*, *Comput. Phys. Commun.*, 195 (2015), pp. 35–48, <https://doi.org/10.1016/j.cpc.2015.04.017>, <https://arxiv.org/abs/1402.0831>.
 - [15] E. ESTRADA, *The Structure of Complex Networks: Theory and Applications*, Oxford University Press, Inc., New York, 2011.
 - [16] E. ESTRADA AND D. J. HIGHAM, *Network properties revealed through matrix functions*, *SIAM Rev.*, 52 (2010), pp. 696–714.
 - [17] A. FROMMER, K. KAHL, S. KRIEG, B. LEDER, AND M. ROTTMAN, *Adaptive Aggregation Based Domain Decomposition Multigrid for the Lattice Wilson Dirac Operator*, *SIAM J. Sci. Comp.*, 36 (2014), pp. A1581–A1608.
 - [18] A. FROMMER, C. SCHIMMEL, AND M. SCHWEITZER, *Analysis of probing techniques for sparse approximation and trace estimation of decaying matrix functions*, *SIAM J. Matrix Anal. Appl.*, <https://arxiv.org/abs/2009.01589>. accepted for publication.
 - [19] A. S. GAMBHIR, A. STATHOPOULOS, AND K. ORGINOS, *Deflation as a method of variance reduction for estimating the trace of a matrix inverse*, *SIAM J. on Sci. Comput.*, 39 (2017), pp. A532–A558, <https://doi.org/10.1137/16M1066361>.
 - [20] M. B. GILES, *Multilevel Monte Carlo methods.*, *Acta Numer.*, 24 (2015), pp. 259–328, <https://doi.org/10.1017/S096249291500001X>.
 - [21] Y. GINOSAR, I. GUTMAN, T. MANSOUR, AND M. SCHORK, *Estrada index and Chebyshev polynomials*, *Chem. Phys. Lett.*, 454 (2008), pp. 145–147.
 - [22] L. GIUSTI, P. HERNANDEZ, M. LAINE, P. WEISZ, AND H. WITTIG, *Low-energy Couplings of QCD from Current Correlators near the Chiral Limit*, *JHEP*, 04 (2004), p. 013, <https://doi.org/10.1088/1126-6708/2004/04/013>, <https://arxiv.org/abs/hep-lat/0402002>.
 - [23] G. H. GOLUB, M. HEATH, AND G. WAHBA, *Generalized cross-validation as a method for choosing a good ridge parameter*, *Technometrics*, 21 (1979), pp. 215–223.
 - [24] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, fourth ed., 2013.
 - [25] G. H. GOLUB AND U. VON MATT, *Generalized cross-validation for large scale problems*, *J. Comput. Graph. Statist.*, 6 (1995), pp. 1–34.
 - [26] I. HAN, D. MALIOUTOV, H. AVRON, AND J. SHIN, *Approximating spectral sums of large-scale matrices using stochastic Chebyshev approximations*, *SIAM J. Sci. Comput.*, 39 (2017), pp. A1558–A1585, <https://doi.org/10.1137/16M1078148>, <https://doi.org/10.1137/16M1078148>.
 - [27] I. HAN, D. MALIOUTOV, AND J. SHIN, *Large-scale log-determinant computation through stochastic Chebyshev expansions*, in *International Conference on Machine Learning*, 2015, pp. 908–917.
 - [28] N. J. HIGHAM, *Functions of Matrices: Theory and Computation*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008.
 - [29] M. F. HUTCHINSON, *A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines*, *Comm. Statist. Simulation Comput.*, 19 (1990), pp. 433–450, <https://doi.org/10.1080/03610919008812864>, <https://doi.org/10.1080/03610919008812864>.
 - [30] J. LAEUCHLI AND A. STATHOPOULOS, *Extending hierarchical probing for computing the trace of matrix inverses*, *SIAM J. Sci. Comput.*, 42 (2020), pp. A1459–A1485.
 - [31] L. LIN, Y. SAAD, AND C. YANG, *Approximating spectral densities of large matrices*, *SIAM Rev.*, 58 (2016), pp. 34–65, <https://doi.org/10.1137/130934283>, <https://doi.org/10.1137/130934283>.
 - [32] Q. LIU, W. WILCOX, AND R. MORGAN, *Polynomial subtraction method for disconnected quark loops*, 2014, <https://arxiv.org/abs/1405.1763>.
 - [33] R. A. MEYER, C. MUSCO, AND D. P. WOODRUFF, *Hutch++: Optimal stochastic trace estimation*, in *Symposium on Simplicity in Algorithms (SOSA)*, SIAM, 2021, pp. 142–155.
 - [34] L. N. OLSON AND J. B. SCHRODER, *PyAMG: Algebraic multigrid solvers in Python v4.0*, 2018, <https://github.com/pyamg/pyamg>. Release 4.0.
 - [35] C. E. RASMUSSEN AND C. K. I. WILLIAMS, *Gaussian Processes for Machine Learning (Adaptive*

- Computation and Machine Learning*), The MIT Press, 2005.
- [36] E. ROMERO, A. STATHOPOULOS, AND K. ORGINOS, *Multigrid deflation for lattice QCD*, J. Comput. Phys., 409 (2020), p. 109356, <https://doi.org/10.1016/j.jcp.2020.109356>, <http://dx.doi.org/10.1016/j.jcp.2020.109356>.
 - [37] F. ROOSTA KHORASANI AND U. ASCHER, *Improved bounds on sample size for implicit matrix trace estimators*, Foundations of Computational Mathematics, 15 (2015), pp. 1187–1212.
 - [38] H. RUE AND L. HELD, *Gaussian Markov Random Fields: Theory And Applications*, CRC press, 2005.
 - [39] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, 2nd edition ed., 2003.
 - [40] B. SAPOVAL, T. GOBRON, AND A. MARGOLINA, *Vibrations of fractal drums*, Phys. Rev. Lett., 67 (1991), pp. 2974–2977.
 - [41] J. SCHWINGER, *Gauge invariance and mass II*, 128 (1962), p. 24252429, <https://doi.org/10.1103/PhysRev.128.2425>.
 - [42] J. SEXTON AND D. WEINGARTEN, *Systematic expansion for full QCD based on the valence approximation*, 1994, <https://arxiv.org/abs/hep-lat/9411029>.
 - [43] A. STATHOPOULOS, J. LAEUCHLI, AND K. ORGINOS, *Hierarchical probing for estimating the trace of the matrix inverse on toroidal lattices*, SIAM J. Sci. Comput., 35 (2013), pp. 299–322.
 - [44] J. M. TANG AND Y. SAAD, *A probing method for computing the diagonal of a matrix inverse*, Numer. Linear Algebra Appl., 19 (2012), pp. 485–501.
 - [45] U. TROTTEBERG, C. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Academic Press, New York, 2000.
 - [46] S. UBARU, J. CHEN, AND Y. SAAD, *Fast estimation of $\text{tr}(f(A))$ via stochastic Lanczos quadrature*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 1075–1099.
 - [47] S. UBARU AND Y. SAAD, *Applications of trace estimation techniques*, in International Conference on High Performance Computing in Science and Engineering, Springer, 2017, pp. 19–33.
 - [48] W. WILCOX, *Noise methods for flavor singlet quantities*, (1999), <https://arxiv.org/abs/hep-lat/9911013v2>.