

2012

TR-2012016: More on the Power of Randomized Matrix Multiplication I

Victor Y. Pan

Guoliang Qian

Follow this and additional works at: http://academicworks.cuny.edu/gc_cs_tr

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Pan, Victor Y. and Qian, Guoliang, "TR-2012016: More on the Power of Randomized Matrix Multiplication I" (2012). *CUNY Academic Works*.
http://academicworks.cuny.edu/gc_cs_tr/376

This Technical Report is brought to you by CUNY Academic Works. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of CUNY Academic Works. For more information, please contact AcademicWorks@gc.cuny.edu.

More on the Power of Randomized Matrix Multiplication I *

Victor Y. Pan^{[1,2],[a]} and Guoliang Qian^{[2],[b]}

^[1] Department of Mathematics and Computer Science
Lehman College of the City University of New York
Bronx, NY 10468 USA

^[2] Ph.D. Programs in Mathematics and Computer Science
The Graduate Center of the City University of New York
New York, NY 10036 USA

^[a] victor.pan@lehman.cuny.edu

<http://comet.lehman.cuny.edu/vpan/>

^[b] gqian@gc.cuny.edu

Abstract

A random matrix is likely to be well conditioned, and motivated by this well known property we employ random matrix multipliers to advance some fundamental matrix computations. This includes numerical stabilization of Gaussian elimination with no pivoting as well as block Gaussian elimination, approximation of the leading and trailing singular spaces of an ill conditioned matrix, associated with its largest and smallest singular values, respectively, and approximation of this matrix by low-rank matrices, with further extensions to Tensor Train approximation and the computation of the numerical rank of a matrix. We formally support the efficiency of the proposed techniques where we employ Gaussian random multipliers, but our extensive tests have consistently produced the same outcome where instead we used sparse and structured random multipliers, defined by much fewer random parameters compared to the number of their entries.

2000 Math. Subject Classification: 15A52, 15A12, 15A06, 65F22, 65F05

Key Words: Random matrices, Random multipliers, GENP, Low-rank approximation, Numerical rank

1 Introduction

It is well known that A random matrix is likely to be well conditioned [D88], [E88], [ES05], [CD05], [SST06], [B11], and motivated by this well known property we apply randomized matrix multiplication to advance some fundamental matrix computations. We stabilize numerically Gaussian elimination with no pivoting as well as block Gaussian elimination, approximate leading and trailing singular spaces of an ill conditioned matrix A , associated with its largest and smallest singular

*Some results of this paper have been presented at the ACM-SIGSAM International Symposium on Symbolic and Algebraic Computation (ISSAC '2011), San Jose, CA, 2011, the 3rd International Conference on Matrix Methods in Mathematics and Applications (MMMA 2011) in Moscow, Russia, June 22-25, 2011, the 7th International Congress on Industrial and Applied Mathematics (ICIAM 2011), in Vancouver, British Columbia, Canada, July 18-22, 2011, the SIAM International Conference on Linear Algebra, in Valencia, Spain, June 18-22, 2012, and the Conference on Structured Linear and Multilinear Algebra Problems (SLA2012), in Leuven, Belgium, September 10-14, 2012

values, respectively, approximate this matrix by low-rank matrices, and compute Tensor Train approximation, the numerical rank of a matrix, and an approximation of a matrix by a structured matrix lying nearby. Our numerical tests are in good accordance with our formal study, except that in the tests all algorithms have fully preserved their power even where we dramatically decreased the number of random parameters involved by using sparse and structured multipliers.

1.1 Numerically safe Gaussian elimination with no pivoting

Hereafter “flop” stands for “arithmetic operation”, by saying “expect” and “likely” we mean “with probability 1 or close to 1”, $\sigma_j(A)$ denotes the j th largest singular value of an $n \times n$ matrix A , and the ratio $\kappa(A) = \sigma_1(A)/\sigma_\rho(A)$ for $\rho = \text{rank}(A)$ denotes its condition number. $\kappa(A) = \|A\| \|A^{-1}\|$ if $\rho = n$, that is if A is a nonsingular matrix. If this number is large in context, then the matrix A is *ill conditioned*, otherwise *well conditioned*. For matrix inversion and solving linear systems of equations the condition number represents the output magnification of input errors,

$$\kappa(A) \approx \frac{\|\text{OUTPUT ERROR}\|}{\|\text{INPUT ERROR}\|}, \tag{1.1}$$

and backward error analysis implies similar magnification of rounding errors [GL96], [H02], [S98].

To avoid dealing with singular or ill conditioned matrices in Gaussian elimination, one incorporates pivoting, that is row or column interchange. *Gaussian elimination with no pivoting* (hereafter we refer to it as *GENP*) can easily fail in numerical computations with rounding errors, except for the cases where the input matrices are strongly well conditioned, that is where all their leading principal square blocks are nonsingular and well conditioned. In particular diagonally dominant as well as positive definite well conditioned matrices have this property. For such matrices, GENP outperforms Gaussian elimination with pivoting [GL96, page 119]. Random matrices are likely to be strongly well conditioned, but we do not solve random linear systems of equations. We can, however randomize linear systems by applying random multipliers and then can apply GENP. We proposed and tested this approach in [PGMQ, Section 12.2] and [PQZa], and our tests consistently showed its efficiency even where we used just circulant or Householder multipliers filled with integers ± 1 and where we limited randomization to the choice of the signs \pm (see our Table 7.1 and [PQZa, Table 2]). Our Corollary 4.1 supports these empirical observations provided that the multipliers are square Gaussian random matrices. Estimation of the condition numbers of structured matrices was stated as a challenge in [SST06]. The paper [PQa] presents some initial advance, but the problem remains largely open.

1.2 Randomized low-rank approximation and beyond

Our Corollary 4.1, supporting randomized GENP, relies on the probabilistic estimates for the ranks and condition numbers of the products $\kappa(GA)$ and $\kappa(AH)$ in terms of $\kappa(A)$ where G and H are Gaussian random matrices (see Theorem 4.1). We also apply the same estimates to support randomized algorithms for the approximation of the leading singular spaces of an ill conditioned matrix A associated with its largest singular values. This can be immediately extended to the approximation of a matrix having a small numerical rank by low-rank matrices. The algorithm is numerically safe, runs at a low computational cost, and has a great number of highly important applications to matrix computations [HMT11]. We point out its further extensions to the approximation of a matrix by a structured matrix lying nearby and to computing Tensor Train approximation and the numerical rank of a matrix. Then again our formal support of these algorithms relies on using Gaussian random multipliers, but our tests show that random Toeplitz multipliers are as effective. This suggests formal and experimental study of various other random structured and sparse multipliers that depend on smaller numbers of random parameters. Note the recent success of Tropp [T11] in this direction.

1.3 Related work

Preconditioning of linear systems of equations is a classical subject [A94], [B02], [G97]. Randomized multiplicative preconditioning for numerical stabilization of GENP was proposed in [PGMQ, Section 12.2] and [PQZa], but with no formal support for this approach. On low-rank approximation we refer the reader to the survey [HMT11]. We cite these and other related works throughout the paper and refer to [PQZb, Section 11] on further bibliography. For a natural extension of our present work, one can combine randomized matrix multiplication with randomized augmentation and additive preprocessing of [PGMQ], [PIMR10], [PQ10], [PQ12], [PQZC], [PQZb], [PY09].

1.4 Organization of the paper and selective reading

In the next section we recall some definitions and basic results. We estimate the condition numbers of Gaussian random matrices in Section 3 and of randomized matrix products in Section 4, where we also comment on numerical stabilization of GENP by means of randomized multiplication. In Sections 5 and 6 we apply randomized matrix multiplication to approximate the leading and trailing singular spaces of a matrix having a small numerical rank, approximate this matrix by a low-rank matrix, and point out applications to tensor decomposition and to approximation by structured matrices. In Section 7 we cover numerical tests, which constitute the contribution of the second author. In Appendix A we estimate the probability that a random matrix has full rank under the uniform probability distribution. In Appendix B we compute the numerical rank of a matrix by using randomization but neither pivoting nor orthogonalization.

2 Some definitions and basic results

We assume computations in the field \mathbb{R} of real numbers.

Hereafter “flop” stands for “arithmetic operation”; “expect” and “likely” mean “with probability 1 or close to 1” (we do not use the concept of the expected value), and the concepts “large”, “small”, “near”, “closely approximate”, “ill conditioned” and “well conditioned” are quantified in the context. Next we recall and extend some customary definitions of matrix computations [GL96], [S98].

2.1 Some basic definitions on matrix computations

$\mathbb{R}^{m \times n}$ is the class of real $m \times n$ matrices $A = (a_{i,j})_{i,j}^{m,n}$.

$(B_1 \mid \dots \mid B_k) = (B_j)_{j=1}^k$ is a $1 \times k$ block matrix with blocks B_1, \dots, B_k . $\text{diag}(B_1, \dots, B_k) = \text{diag}(B_j)_{j=1}^k$ is a $k \times k$ block diagonal matrix with diagonal blocks B_1, \dots, B_k .

\mathbf{e}_i is the i th coordinate vector of dimension n for $i = 1, \dots, n$. These vectors define the identity matrix $I_n = (\mathbf{e}_1 \mid \dots \mid \mathbf{e}_n)$ of size $n \times n$. $O_{k,l}$ is the $k \times l$ matrix filled with zeros. We write I and O where the size of a matrix is not important or is defined by context.

A^T is the transpose of a matrix A .

2.2 Range, rank, and generic rank profile

$\mathcal{R}(A)$ denotes the range of an $m \times n$ matrix A , that is the linear space $\{\mathbf{z} : \mathbf{z} = A\mathbf{x}\}$ generated by its columns. $\text{rank}(A) = \dim \mathcal{R}(A)$ denotes its rank. $A_k^{(k)}$ denotes the leading, that is northwestern $k \times k$ block submatrix of a matrix A . A matrix of a rank ρ has *generic rank profile* if all its leading $i \times i$ blocks are nonsingular for $i = 1, \dots, \rho$. If such matrix is nonsingular itself, then it is called *strongly nonsingular*.

Fact 2.1. *The set \mathbb{M} of $m \times n$ matrices of rank ρ is an algebraic variety of dimension $(m + n - \rho)\rho$.*

Proof. Let M be an $m \times n$ matrix of a rank ρ with a nonsingular leading $\rho \times \rho$ block M_{00} and write $M = \begin{pmatrix} M_{00} & M_{01} \\ M_{10} & M_{11} \end{pmatrix}$. Then the $(m - \rho) \times (n - \rho)$ Schur complement $M_{11} - M_{10}M_{00}^{-1}M_{01}$ must vanish, which imposes $(m - \rho)(n - \rho)$ algebraic equations on the entries of M . Similar argument

can be applied where any $\rho \times \rho$ submatrix of the matrix M (among $\binom{m}{\rho} \binom{n}{\rho}$ such submatrices) is nonsingular. Therefore $\dim \mathbb{M} = mn - (m - \rho)(n - \rho) = (m + n - \rho)\rho$. \square

2.3 Orthogonal, Toeplitz and circulant matrices

A real matrix Q is called *orthogonal* if $Q^T Q = I$ or $Q Q^T = I$. In Section 7 we write $Q(A)$ to denote a unique orthogonal matrix specified by the following result.

Fact 2.2. [GL96, Theorem 5.2.2]. *QR factorization* $A = QR$ of a matrix A having full column rank into the product of an orthogonal matrix $Q = Q(A)$ and an upper triangular matrix $R = R(A)$ is unique provided that the factor R is a square matrix with positive diagonal entries.

A *Toeplitz* $m \times n$ matrix $T_{m,n} = (t_{i-j})_{i,j=1}^{m,n}$ is defined by its first row $(t_{-h})_{h=0}^{n-1}$ and the subvector $(t_h)_{h=1}^{n-1}$ of its first column vector. Circulant matrices are the subclass of Toeplitz matrices where $t_g = t_h$ if $|g - h| = n$.

Theorem 2.1. $O((m+n) \log(m+n))$ flops suffice to multiply an $m \times n$ Toeplitz matrix by a vector.

2.4 Norms, SVD, generalized inverse, and singular spaces

$\|A\|_h$ is the h -norm and $\|A\|_F = \sqrt{\sum_{i,j=1}^{m,n} |a_{i,j}|^2}$ is the Frobenius norm of a matrix $A = (a_{i,j})_{i,j=1}^{m,n}$. We write $\|A\| = \|A\|_2$ and $\|\mathbf{v}\| = \sqrt{\mathbf{v}^T \mathbf{v}} = \|\mathbf{v}\|_2$ and recall from [GL96, Section 2.3.2 and Corollary 2.3.2] that

$$\max_{i,j=1}^{m,n} |a_{i,j}| \leq \|A\| = \|A^T\| \leq \sqrt{mn} \max_{i,j=1}^{m,n} |a_{i,j}|,$$

$$\frac{1}{\sqrt{m}} \|A\|_1 \leq \|A\| \leq \sqrt{n} \|A\|_1, \quad \|A\|_1 = \|A^T\|_\infty, \quad \|A\|^2 \leq \|A\|_1 \|A\|_\infty, \quad (2.1)$$

$$\|A\| \leq \|A\|_F \leq \sqrt{n} \|A\|, \quad (2.2)$$

$$\|AB\|_h \leq \|A\|_h \|B\|_h \text{ for } h = 1, 2, \infty \text{ and any matrix product } AB. \quad (2.3)$$

Define an *SVD* or *full SVD* of an $m \times n$ matrix A of a rank ρ as follows,

$$A = S_A \Sigma_A T_A^T. \quad (2.4)$$

Here $S_A S_A^T = S_A^T S_A = I_m$, $T_A T_A^T = T_A^T T_A = I_n$, $\Sigma_A = \text{diag}(\widehat{\Sigma}_A, O_{m-\rho, n-\rho})$, $\widehat{\Sigma}_A = \text{diag}(\sigma_j(A))_{j=1}^\rho$, $\sigma_j = \sigma_j(A) = \sigma_j(A^T)$ is the j th largest singular value of a matrix A for $j = 1, \dots, \rho$, and we write $\sigma_j = 0$ for $j > \rho$. These values have the minimax property

$$\sigma_j = \max_{\dim(\mathbb{S})=j} \min_{\mathbf{x} \in \mathbb{S}, \|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\|, \quad j = 1, \dots, \rho, \quad (2.5)$$

where \mathbb{S} denotes linear spaces [GL96, Theorem 8.6.1]. Consequently $\sigma_\rho > 0$, $\sigma_1 = \max_{\|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\| = \|A\|$.

Fact 2.3. *If A_0 is a submatrix of a matrix A , then $\sigma_j(A) \geq \sigma_j(A_0)$ for all j .*

Proof. [GL96, Corollary 8.6.3] implies the claimed bound where A_0 is any block of columns of the matrix A . Transposition of a matrix and permutations of its rows and columns do not change singular values, and thus we can extend the bounds to all submatrices A_0 . \square

$A^+ = T_A \text{diag}(\widehat{\Sigma}_A^{-1}, O_{n-\rho, m-\rho}) S_A^T$ is the Moore–Penrose pseudo-inverse of the matrix A of (2.4), and

$$\|A^+\| = 1/\sigma_\rho(A) \quad (2.6)$$

for a matrix A of a rank ρ . A^{+T} stands for $(A^+)^T = (A^T)^+$, and A^{-T} stands for $(A^{-1})^T = (A^T)^{-1}$.

In Sections 5–6 we use the following definitions. For every integer k in the range $1 \leq k < \text{rank}(A)$ define the partition $S_A = (S_{k,A} \mid S_{A,m-k})$ and $T_A = (T_{k,A} \mid T_{A,n-k})$ where the submatrices $S_{k,A}$

and $T_{k,A}$ are formed by the first k columns of the matrices S_A and T_A , respectively. Write $\Sigma_{k,A} = \text{diag}(\sigma_j(A))_{j=1}^k$, $\mathbb{S}_{k,A} = \mathcal{R}(S_{k,A})$ and $\mathbb{T}_{k,A} = \mathcal{R}(T_{k,A})$. If $\sigma_k > \sigma_{k+1}$, then $\mathbb{S}_{k,A}$ and $\mathbb{T}_{k,A}$ are the left and right *leading singular spaces*, respectively, associated with the k largest singular values of the matrix A , whereas their orthogonal complements $\mathbb{S}_{A,m-k} = \mathcal{R}(S_{A,m-k})$ and $\mathbb{T}_{A,n-k} = \mathcal{R}(T_{A,n-k})$ are the left and right *trailing singular spaces*, respectively, associated with the other singular values of A . The pairs of subscripts $\{k, A\}$ versus $\{A, m - k\}$ and $\{A, n - k\}$ mark the leading versus trailing singular spaces. The left singular spaces of A are the right singular spaces of A^T and vice versa. All matrix bases for the singular spaces $\mathbb{S}_{k,A}$ and $\mathbb{T}_{k,A}$ are given by matrices $S_{k,A}X$ and $T_{k,A}Y$, respectively, for nonsingular $k \times k$ matrices X and Y . Orthogonal matrices X and Y define orthogonal matrix bases for these spaces. B is an *approximate matrix basis* for a space \mathbb{S} within a relative error norm bound τ if there exists a matrix E such that $B + E$ is a matrix basis for this space \mathbb{S} and if $\|E\| \leq \tau\|B\|$.

2.5 Condition number, numerical rank and generic conditioning profile

$\kappa(A) = \frac{\sigma_1(A)}{\sigma_\rho(A)} = \|A\| \|A^+\|$ is the condition number of an $m \times n$ matrix A of a rank ρ . Such matrix is *ill conditioned* if $\sigma_1(A) \gg \sigma_\rho(A)$ and is *well conditioned* otherwise. See [D83], [GL96, Sections 2.3.2, 2.3.3, 3.5.4, 12.5], [H02, Chapter 15], [KL94], [S98, Section 5.3], on the estimation of matrix norms and condition numbers.

An $m \times n$ matrix A has *numerical rank*, denoted $\text{nrnk}(A)$ and not exceeding $\text{rank}(A)$, if the ratios $\sigma_j(A)/\|A\|$ are small for $j > \text{nrnk}(A)$ but not for $j \leq \text{nrnk}(A)$.

Remark 2.1. *One can specify the adjective “small” above as “smaller than a fixed positive tolerance”. The choice of the tolerance can be a challenge, e.g., for the matrix $\text{diag}(1.1^{-j})_{j=0}^{999}$.*

If a well conditioned $m \times n$ matrix A has a rank $\rho < l = \min\{m, n\}$, then almost all its close neighbours have full rank l (see Section 3.2), and all of them have numerical rank ρ . Conversely, suppose a matrix A has a positive numerical rank $\rho = \text{nrnk}(A)$ and *truncate its SVD* by setting to 0 all its singular values, except for the ρ largest ones. Then the resulting matrix $A - E$ is well conditioned and has rank ρ and $\|E\| = \sigma_{\rho+1}(A)$, and so $A - E$ is a rank- ρ approximation to the matrix A within the error norm bound $\sigma_{\rho+1}(A)$. At a lower computational cost we can obtain rank- ρ approximations of the matrix A from its rank-revealing factorizations [GE96], [HP92], [P00a], and we further decrease the computational cost by applying randomized algorithms in Section 5.

An $m \times n$ matrix has *generic conditioning profile* (cf. the end of Section 2.2) if it has a numerical rank ρ and if its leading $i \times i$ blocks are nonsingular and well conditioned for $i = 1, \dots, \rho$. If such matrix has full rank (that is if $\rho = \min\{m, n\}$) and if it is well conditioned itself, then we call it *strongly well conditioned*. The following theorem shows that GENP and block Gaussian elimination applied to a strongly well conditioned matrix are numerically safe.

Theorem 2.2. *Cf. [PQZa, Theorem 5.1]. Assume GENP or block Gaussian elimination applied to an $n \times n$ matrix A and write $N = \|A\|$ and $N_- = \max_{j=1}^n \|(A_j^{(j)})^{-1}\|$. Then the absolute values of all pivot elements of GENP and the norms of all pivot blocks of block Gaussian elimination do not exceed $N + N_-N^2$, whereas the absolute values of the reciprocals of these elements and the norms of the inverses of the blocks do not exceed N_- .*

3 Ranks and conditioning of Gaussian random matrices

3.1 Random variables and Gaussian random matrices

Definition 3.1. $F_\gamma(y) = \text{Probability}\{\gamma \leq y\}$ (for a real random variable γ) is the cumulative distribution function (cdf) of γ evaluated at y . $F_{g(\mu, \sigma)}(y) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^y \exp(-\frac{(x-\mu)^2}{2\sigma^2}) dx$ for a Gaussian random variable $g(\mu, \sigma)$ with a mean μ and a positive variance σ^2 , and so

$$\mu - 4\sigma \leq y \leq \mu + 4\sigma \text{ with a probability near 1.} \quad (3.1)$$

Definition 3.2. A matrix (or a vector) is a Gaussian random matrix (or vector) with a mean μ and a positive variance σ^2 if it is filled with independent identically distributed Gaussian random variables, all having the mean μ and variance σ^2 . $\mathcal{G}_{\mu,\sigma}^{m \times n}$ is the set of such Gaussian random $m \times n$ matrices, which are standard for $\mu = 0$ and $\sigma^2 = 1$. By restricting this set to $m \times n$ Toeplitz matrices where only the $m + n - 1$ entries of the first row and column are independent we obtain the set of $\mathcal{T}_{\mu,\sigma}^{m \times n}$ Gaussian random Toeplitz matrices. Likewise we obtain the set $\mathcal{Z}_{\mu,\sigma}^{n \times n}$ of Gaussian random circulant matrices, where only n entries of the first row are independent.

3.2 Nondegeneration of Gaussian random matrices

The total degree of a multivariate monomial is the sum of its degrees in all its variables. The total degree of a polynomial is the maximal total degree of its monomials.

Lemma 3.1. [DL78], [S80], [Z79]. For a set Δ of a cardinality $|\Delta|$ in any fixed ring let a polynomial in m variables have a total degree d and let it not vanish identically on this set. Then the polynomial vanishes in at most $d|\Delta|^{m-1}$ points.

We assume that Gaussian random variables range over infinite sets Δ , usually over the real line or its interval. Then the lemma implies that a nonzero polynomial vanishes with probability 0. Consequently a Gaussian random general, Toeplitz or circulant matrix has generic rank profile with probability 1 because the determinant of any its block is a polynomials in the entries. Likewise Gaussian random general, Toeplitz and circulant matrices have generic rank profile with probability 1. Hereafter, wherever this causes no confusion, we assume by default that *Gaussian random general, Toeplitz and circulant matrices have generic rank profile*. This property can be readily extended to the products and various functions of general, sparse and structured Gaussian random matrices. Similar properties hold with probability near 1 where the random variables are sampled under the uniform probability distribution from a finite set of a large cardinality (see the Appendix).

3.3 Extremal singular values of Gaussian random matrices

Besides having full rank with probability 1, Gaussian random matrices in Definition 3.2 are likely to be well conditioned [D88], [E88], [ES05], [CD05], [B11], and even the sum $M + A$ for $M \in \mathbb{R}^{m \times n}$ and $A \in \mathcal{G}_{\mu,\sigma}^{m \times n}$ is likely to be well conditioned unless the ratio $\sigma/||M||$ is small or large [SST06].

The following theorem states an upper bound proportional to y on the cdf $F_{1/||A^+||}(y)$, that is on the probability that the smallest positive singular value $1/||A^+|| = \sigma_l(A)$ of a Gaussian random matrix A is less than a nonnegative scalar y (cf. (2.6)) and consequently on the probability that the norm $||A^+||$ exceeds a positive scalar x . The stated bound still holds if we replace the matrix A by $A - B$ for any fixed matrix B , and for $B = O_{m,n}$ the bounds can be strengthened by a factor $y^{|m-n|}$ [ES05], [CD05].

Theorem 3.1. Suppose $A \in \mathcal{G}_{\mu,\sigma}^{m \times n}$, $B \in \mathbb{R}^{m \times n}$, $l = \min\{m, n\}$, $x > 0$, and $y \geq 0$. Then $F_{\sigma_l(A-B)}(y) \leq 2.35 \sqrt{l}y/\sigma$, that is $\text{Probability}\{||(A - B)^+|| \geq 2.35x\sqrt{l}/\sigma\} \leq 1/x$.

Proof. For $m = n$ this is [SST06, Theorem 3.3]. Apply Fact 2.3 to extend it to any pair $\{m, n\}$. \square

The following two theorems supply lower bounds $F_{||A||}(z)$ and $F_{\kappa(A)}(y)$ on the probabilities that $||A|| \leq z$ and $\kappa(A) \leq y$ for two scalars y and z , respectively, and a Gaussian random matrix A . We do not use the second theorem, but state it for the sake of completeness and only for square $n \times n$ matrices A . The theorems imply that the functions $1 - F_{||A||}(z)$ and $1 - F_{\kappa(A)}(y)$ decay as $z \rightarrow \infty$ and $y \rightarrow \infty$, respectively, and that the decays are exponential in $-z^2$ and proportional to $\sqrt{\log y}/y$, respectively. For small values $y\sigma$ and a fixed n the lower bound of Theorem 3.3 becomes negative, in which case the theorem becomes trivial. Unlike Theorem 3.1, in both theorems we assume that $\mu = 0$.

Theorem 3.2. [DS01, Theorem II.7]. Suppose $A \in \mathcal{G}_{0,\sigma}^{m \times n}$, $h = \max\{m, n\}$ and $z \geq 2\sigma\sqrt{h}$. Then $F_{||A||}(z) \geq 1 - \exp(-(z - 2\sigma\sqrt{h})^2/(2\sigma^2))$, and so the norm $||A||$ is likely to have order $\sigma\sqrt{h}$.

Theorem 3.3. [SST06, Theorem 3.1]. Suppose $0 < \sigma \leq 1$, $y \geq 1$, $A \in \mathcal{G}_{0,\sigma}^{n \times n}$. Then the matrix A has full rank with probability 1 and $F_{\kappa(A)}(y) \geq 1 - (14.1 + 4.7\sqrt{(2 \ln y)/n})n/(y\sigma)$.

Proof. See [SST06, the proof of Lemma 3.2]. \square

4 Condition numbers of randomized matrix products and generic preconditioning

Next we deduce probabilistic lower bounds on the smallest singular values of the products of fixed and random matrices. We begin with three lemmas. The first of them is obvious, the second easily follows from minimax property (2.5).

Lemma 4.1. $\sigma_j(SM) = \sigma_j(MT) = \sigma_j(M)$ for all j if S and T are square orthogonal matrices.

Lemma 4.2. Suppose $\Sigma = \text{diag}(\sigma_i)_{i=1}^n$, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$, $G \in \mathbb{R}^{r \times n}$, $H \in \mathbb{R}^{n \times r}$. Then $\sigma_j(G\Sigma) \geq \sigma_j(G)\sigma_n$, $\sigma_j(\Sigma H) \geq \sigma_j(H)\sigma_n$ for all j . If also $\sigma_n > 0$, then $\text{rank}(G\Sigma) = \text{rank}(G)$, $\text{rank}(\Sigma H) = \text{rank}(H)$.

Lemma 4.3. [SST06, Proposition 2.2]. Suppose $H \in \mathcal{G}_{\mu,\sigma}^{m \times n}$, $SS^T = S^T S = I_m$, $TT^T = T^T T = I_n$. Then $SH \in \mathcal{G}_{\mu,\sigma}^{m \times n}$ and $HT \in \mathcal{G}_{\mu,\sigma}^{m \times n}$.

The following theorem implies that multiplication by standard Gaussian random matrix is unlikely to decrease the smallest positive singular value of a matrix dramatically, even though $UV = O$ for some pairs of rectangular orthogonal matrices U and V .

Theorem 4.1. Suppose $G' \in \mathcal{G}_{\mu,\sigma}^{r \times m}$, $H' \in \mathcal{G}_{\mu,\sigma}^{n \times r}$, $M \in \mathbb{R}^{m \times n}$, $G = G' + U$, $H = H' + V$ for some matrices U and V , $r(M) = \text{rank}(M)$, $x > 0$ and $y \geq 0$. Then $F_{1/\|(GM)^+\|}(y) \leq F(y, M, \sigma)$ and $F_{1/\|(MH)^+\|}(y) \leq F(y, M, \sigma)$ for $F(y, M, \sigma) = 2.35y\sqrt{\hat{r}}\|M^+\|/\sigma$ and $\hat{r} = \min\{r, r(M)\}$, that is $\text{Probability}\{\|P^+\| \geq 2.35x\sqrt{\hat{r}}\|M^+\|/\sigma\} \leq 1/x$ for $P = GM$ and $P = MH$.

Proof. With probability 1, the matrix MH has rank \hat{r} because $H \in \mathcal{G}_{\mu,\sigma}^{n \times r}$. So (cf. (2.6))

$$F_{1/\|(MH)^+\|}(y) = F_{\sigma_{\hat{r}}(MH)}(y). \quad (4.1)$$

Let $M = S_M \Sigma_M T_M^T$ be full SVD where $\Sigma_M = \text{diag}(\hat{\Sigma}_M, O) = \Sigma_M \text{diag}(I_{r(M)}, O)$ and $\hat{\Sigma}_M = \text{diag}(\sigma_j(M))_{j=1}^{r(M)}$ is a nonsingular diagonal matrix. We have $MH = S_M \Sigma_M T_M^T H$, and so $\sigma_j(MH) = \sigma_j(\Sigma_M T_M^T H)$ for all j by virtue of Lemma 4.1, because S_M is a square orthogonal matrix. Write $H_{r(M)} = (I_{r(M)} \mid O) T_M^T H$ and observe that $\sigma_j(\Sigma_M T_M^T H) = \sigma_j(\hat{\Sigma}_M H_{r(M)})$ and consequently

$$\sigma_j(MH) = \sigma_j(\hat{\Sigma}_M H_{r(M)}) \text{ for all } j. \quad (4.2)$$

Combine equation (4.2) for $j = \hat{r}$ with Lemma 4.2 for the pair (Σ, H) replaced by $(\hat{\Sigma}_M, H_{r(M)})$ and obtain that $\sigma_{\hat{r}}(MH) \geq \sigma_{r(M)}(M)\sigma_{\hat{r}}(H_{r(M)}) = \sigma_{\hat{r}}(H_{r(M)})/\|M^+\|$. We have $T_M^T H' \in \mathcal{G}_{\mu,\sigma}^{n \times r}$ by virtue of Lemma 4.3, because T_M is a square orthogonal matrix; consequently $H_{r(M)} = H'_{r(M)} + B$ for $H'_{r(M)} \in \mathcal{G}_{\mu,\sigma}^{r(M) \times r}$ and some matrix B . Therefore we can apply Theorem 3.1 for $A = H'_{r(M)}$ and obtain the bound of Theorem 4.1 on $F_{1/\|(MH)^+\|}(y)$. One can similarly deduce the bound on $F_{1/\|(GM)^+\|}(y)$ or can just apply the above bound on $F_{1/\|(MH)^+\|}(y)$ for $H = G^T$ and M replaced by M^T and then recall that $(M^T G^T)^T = GM$. \square

By combining (2.3) with Theorems 3.2 (for $B = O$) and 4.1 we can probabilistically bound the condition numbers of randomized products GM and MH . The following corollary extends the bound of Theorem 4.1 for a randomized matrix product to the bounds for its blocks.

Corollary 4.1. *Suppose j, k, m, n, q and s are integers, $1 \leq j \leq q, 1 \leq k \leq s, M \in \mathbb{R}^{m \times n}, \sigma > 0, G \in \mathcal{G}_{\mu, \sigma}^{q \times m}, H \in \mathcal{G}_{\mu, \sigma}^{n \times s}, \text{rank}(M_j) = j$ for $M_j = M \begin{pmatrix} I_j \\ O_{n-j, j} \end{pmatrix}, \text{rank}(M^{(k)}) = k$ for $M^{(k)} = (I_k \mid O_{k, m-k})M$, and $y \geq 0$. Then (i) with probability 1 the matrix GM (resp. MH) has full rank if $\text{rank}(M) \geq q$ (resp. if $\text{rank}(M) \geq s$). Furthermore (ii) $F_{1/\|((GM)_j^{(j)})^+\|}(y) \leq 2.35y\sqrt{j}\|M_j^+\|/\sigma$ if $\text{rank}(M) \geq j, F_{1/\|((MH)_k^{(k)})^+\|}(y) \leq 2.35y\sqrt{k}\|(M^{(k)})^+\|/\sigma$ if $\text{rank}(M) \geq k$.*

Proof. We immediately verify part (i) by applying the techniques of Section 3.2. To prove part (ii) apply Theorem 4.1 replacing G by $(I_j \mid O_{j, q-j})G$ and replacing M by $M \begin{pmatrix} I_j \\ O_{n-j, j} \end{pmatrix}$. For every k apply Theorem 4.1 replacing M by $(I_k \mid O_{k, m-k})M$ and replacing H by $H \begin{pmatrix} I_k \\ O_{s-k, k} \end{pmatrix}$. \square

Corollary 4.1 can be immediately extended to any block of the matrices GM and MH , but we single out the leading blocks because applications of GENP and block Gaussian elimination are numerically safe where these blocks are nonsingular and well conditioned. We have empirical evidence that such applications are numerically safe even where we use circulant multipliers G and H filled with ± 1 and where randomization is restricted to choosing the signs \pm (see Tables 7.1 and 7.3). The paper [T11] provides some formal support for using some other structured randomized multipliers.

5 Approximate bases for singular spaces, low-rank approximation, and the computation of numerical rank

5.1 Randomized low-rank approximation: an outline and an extension to approximation by structured matrices

Suppose we seek a rank- ρ approximation to a matrix A that has a numerical rank ρ . We can solve this problem by computing the SVD of the matrix A or its rank-revealing factorization [GE96], [HP92], [P00a], but in this section we cover alternative numerically stable and noncostly solutions based on randomized matrix multiplication. As by-product we obtain approximate matrix bases for the left or right leading singular space $\mathbb{T}_{\rho, A}$ and $\mathbb{S}_{\rho, A}$.

Let us supply further details. Our next theorem expresses a rank- ρ approximation to a matrix A through a matrix basis for any of the two leading singular spaces $\mathbb{T}_{\rho, A}$ any $\mathbb{S}_{\rho, A}$. Theorem 5.2 of Section 5.3 supports randomized computation of such an approximate basis for the space $\mathbb{T}_{\rho, A}$ from the product $A^T G$ for $G \in \mathcal{G}_{0,1}^{m \times \rho+}$. The paper [T11] formally supports this algorithm for a special class of random structured multipliers G , and our tests consistently show such support where $G \in \mathcal{T}_{0,1}^{m \times \rho+}$ (see Tables 7.2 and 7.3). We conjecture that the same is true for various other classes of sparse and structured multipliers G , defined by much fewer random parameters compared to the number of the entries. We specify a low-rank approximation algorithm in Section 5.4, which has important applications to matrix computations, many listed in [HMT11]. For a natural extension assume a matrix W having a possibly unknown numerical displacement rank d , that is lying near some matrices with a displacement rank d (see the definitions in [KKM79], [BM01], [P01]). We can compute one of these displacements as a rank- d approximation to the displacement of the matrix W , and then immediately recover a structured matrix approximating the matrix W .

5.2 Low-rank approximation via the basis of a leading singular space

Next we prove that both orthogonal and nonorthogonal projections of a matrix A onto its leading singular spaces $\mathbb{T}_{\rho, A}$ and $\mathbb{S}_{\rho, A}$ approximate the matrix within the error norm $\sigma_{\rho+1}(A)$.

Theorem 5.1. *Suppose A is an $m \times n$ matrix, $S_A \Sigma_A T_A^T$ is its SVD of (2.4), ρ is a positive integer, $\rho \leq l = \min\{m, n\}$, and T and S are matrix bases for the spaces $\mathbb{T}_{\rho, A}$ and $\mathbb{S}_{\rho, A}$, respectively. Then*

$$\|A - AT(T^T T)^{-1} T^T\| = \|A - S(S^T S)^{-1} S^T A\| = \sigma_{\rho+1}(A). \quad (5.1)$$

For orthogonal matrices T and S we have $T^T T = S^T S = I_\rho$ and

$$\|A - ATT^T\| = \|A - SS^T A\| = \sigma_{\rho+1}(A). \quad (5.2)$$

Proof. Write $P = T_{\rho,A} T_{\rho,A}^T$ and $r = n - \rho$, observe that $T_A^T T_{\rho,A} = \begin{pmatrix} I_\rho \\ O_{r,\rho} \end{pmatrix}$, and obtain $AP = S_A \Sigma_A T_A^T T_{\rho,A} T_{\rho,A}^T = S_A \Sigma_A \begin{pmatrix} T_{\rho,A}^T \\ O_{r,\rho} \end{pmatrix}$, whereas $A = S_A \Sigma_A \begin{pmatrix} T_{\rho,A}^T \\ T_{A,r}^T \end{pmatrix}$. Hence $A - AP = S_A \Sigma_A \begin{pmatrix} O_{\rho,n} \\ T_{A,r}^T \end{pmatrix} = S_A \text{diag}(O_{\rho,\rho}, \Sigma_\rho) T_A^T$, where Σ_ρ is the $(m - \rho) \times (n - \rho)$ diagonal matrix with the diagonal entries $\sigma_{\rho+1}, \dots, \sigma_l$. Thus $\|A - AP\| = \|\Sigma_\rho\| = \sigma_{\rho+1}(A)$ because S_A and T_A are square orthogonal matrices. This proves the estimates (5.1) and (5.2) for $T = T_{\rho,A}$. Let us extend them to any matrix basis T for the space $\mathbb{T}_{\rho,A}$, that is for $T = T_{\rho,A} U$ where U is a nonsingular matrix. Recall that $T_{\rho,A}^T T_{\rho,A} = I_\rho$, and obtain successively $(U^T T_{\rho,A}^T T_{\rho,A} U)^{-1} = U^{-1} U^{-T}$, $U (U^T T_{\rho,A}^T T_{\rho,A} U)^{-1} U^T = I_\rho$, and $T (T^T T)^{-1} T^T = T_{\rho,A} U (U^T T_{\rho,A}^T T_{\rho,A} U)^{-1} U^T T_{\rho,A}^T = T_{\rho,A} T_{\rho,A}^T$, implying the desired extension. Apply the proof to the transpose A^T to extend it to matrix bases S for the space $\mathbb{S}_{\rho,A}$. \square

5.3 A basis of a leading singular space via randomized products

The following theorem supports randomized approximation of matrix bases for the leading singular spaces $\mathbb{T}_{\rho,A}$ and $\mathbb{S}_{\rho,A}$ of a matrix A .

Theorem 5.2. *Suppose $A \in \mathbb{R}^{m \times n}$, $H \in \mathcal{G}_{0,1}^{n \times \rho}$, and $G \in \mathcal{G}_{0,1}^{m \times \rho}$ and write $S = AH$ and $T = A^T G$. Then (i) $\text{rank}(T) = \text{rank}(S) = \min\{\rho, \text{rank}(A)\}$ with probability 1 and (ii) $\text{nrank}(T) = \text{nrank}(S) = \min\{\rho, \text{nrank}(A)\}$ with probability close to 1. (iii) Furthermore with probability close to 1 we have*

$$S + \Delta = S_{\rho,A} U \quad \text{and} \quad T + \Delta' = T_{\rho,A} V \quad (5.3)$$

for two matrices Δ and Δ' having norms of order $\sigma_{\rho+1}(A)$ and for two nonsingular matrices U and V having condition numbers of at most order $\|A\|/(\sigma_\rho(A)\sqrt{\rho})$.

Proof. We prove the claims about the matrix S , then apply them to the transpose A^T to extend to the matrix T . The techniques of Section 3.2 support part (i). By truncating the SVD of the matrix A to the level of its numerical rank ρ_- we obtain a well conditioned matrix having both rank and numerical rank ρ_- . Then we deduce part (ii) from Theorem 4.1.

Proving part (iii) we assume w.l.o.g. that $\text{rank}(A) \geq \rho$. (Otherwise infinitesimal perturbation of the matrix A could yield this bound.) Write SVD $A = S_A \Sigma_A T_A^T$, $\Sigma_{\rho,A} = \text{diag}(\sigma_j(A))_{j=1}^\rho$, $U = \Sigma_{\rho,A} T_{\rho,A}^T H$, and $A_\rho = S_{\rho,A} U = S_A \text{diag}(\Sigma_{\rho,A}, O_{m-\rho, n-\rho}) T_A^T = S_{\rho,A} \Sigma_{\rho,A} T_{\rho,A}^T$. Note that $\|\Sigma_A - \text{diag}(\Sigma_{\rho,A}, O_{m-\rho, n-\rho})\| = \sigma_{\rho+1}(A)$. Consequently $\|A - A_\rho\| = \sigma_{\rho+1}(A)$, $AH = A_\rho H + \Delta$, where $\|\Delta\| \leq \sigma_{\rho+1}(A) \|H\|$, and the norm $\|H\|$ is likely to be bounded from above and below by two positive constants (see Theorem 3.2). This implies (5.3).

With probability 1 the $\rho \times \rho$ matrices $U = \Sigma_{\rho,A} B$ and $B = T_{\rho,A} H$ are nonsingular (see Section 3.2). Next we assume that they are nonsingular and estimate $\kappa(U)$. Clearly $\|U\| \leq \|\Sigma_{\rho,A}\| \|T_{\rho,A}^T\| \|H\|$ where $\|\Sigma_{\rho,A}\| = \|A\|$, $\|T_{\rho,A}^T\| = 1$. Therefore $\|U\| \leq \|A\| \|H\|$, which is likely to have order $\|A\|$.

Furthermore we have $\|U^{-1}\| \leq \|\Sigma_{\rho,A}^{-1}\| \|B^{-1}\| = \|B^{-1}\|/\sigma_\rho(A)$. Apply Theorem 4.1 where $M = T_{\rho,A}^T \hat{r}$, $\hat{r} = \rho$ and $\sigma_{r(M)}(M) = \sigma = 1$ and obtain that the norm $\|B^{-1}\|$ is likely to have at most order $1/\sqrt{\rho}$. Therefore with probability close to 1, the norm $\|U^{-1}\| = \|B^{-1}\|/\sigma_\rho(A)$ has at most order $1/(\sigma_\rho(A)\sqrt{\rho})$, and then $\kappa(U) = \|U\| \|U^{-1}\|$ has at most order $\|A\|/(\sigma_\rho(A)\sqrt{\rho})$. \square

5.4 A prototype algorithm for low-rank approximation

Together Theorems 5.1 and 5.2 imply correctness of the following prototype algorithm where we assume that the input matrix has an unknown numerical rank and we know its upper bound. The algorithm employs approximation of a leading singular space of the input matrix.

Proto-Algorithm 5.1. Rank- ρ approximation of a matrix (cf. [HMT11, Section 10.3]).

INPUT: A matrix $A \in \mathbb{R}^{m \times n}$ having an unknown numerical rank ρ , an integer $\rho_+ \geq \rho$, and two tolerances τ and τ' of order $\sigma_{\rho+1}(A)/\|A\|$. (We can choose τ at Stage 2 based on rank revealing factorization of an auxiliary $n \times \rho_+$ matrix. The computation of this factorization is noncostly where ρ is small. We can choose τ' at Stage 3 based on the required output accuracy, and can adjust both tolerances if the algorithm fails to produce a satisfactory output.)

OUTPUT: FAILURE (with a low probability) or an integer ρ and two matrices $T \in \mathbb{R}^{n \times \rho}$ and $\hat{A}_\rho \in \mathbb{R}^{m \times n}$, both having ranks at most ρ and such that $\|\hat{A}_\rho - A\| \leq \tau'\|A\|$ and T satisfies bound (5.3) of Theorem 5.2 for $\|\Delta'\| \leq \tau\|A\|$.

COMPUTATIONS:

1. Compute the $n \times \rho_+$ matrix $T' = A^T G$ for $G \in \mathcal{G}_{0,1}^{m \times \rho_+}$.
2. Compute a rank revealing factorization of the matrix T' and choose the minimal integer s and an $n \times s$ matrix T such that $\|T' - (T \mid O_{n, \rho_+ - s})\| \leq \tau\|A\|$.
3. Compute the matrix $\hat{A}_s = AT(T^T T)^{-1}T^T$. Output $\rho = s$, T and \hat{A}_ρ and stop if $\|\hat{A}_\rho - A\| \leq \tau'\|A\|$. Otherwise output FAILURE and stop.

Assume a proper choice of both tolerances τ and τ' . Then by virtue of Theorem 5.2, we can expect that at Stage 2 we obtain $s = \rho$ and an approximate matrix basis T for the singular space $\mathbb{T}_{\rho, A}$ (within an error norm of at most order $\sigma_{\rho+1}(A)$). If so, Stage 3 outputs FAILURE with a probability near 0, by virtue of Theorems 5.1, and in the case of FAILURE we can reapply the algorithm for new values of random parameters or for the adjusted tolerance values τ and τ' . At Stage 2 we have $s \leq \rho$ because $\text{nrnk}(A^T G) \leq \text{nrnk}(A) = \rho$, whereas for a sufficiently small tolerance τ' the bound $\|\hat{A}_\rho - A\| \leq \tau'\|A\|$ at Stage 3 implies that $s \geq \text{nrnk}(A)$. These observations enable us to certify correctness of the outputs ρ , T , and \hat{A}_ρ of the algorithm.

We can similarly approximate the matrix A by a rank- ρ matrix $S(S^T S)^{-1}S^T A$, by first computing the matrix $S' = AH$ for $H \in \mathcal{G}_{0,1}^{n \times \rho_+}$, then computing its rank revealing factorization, which is expected to define an approximate matrix basis S for the space $\mathbb{S}_{\rho, A}$, and finally applying Theorem 5.1, to approximate the matrix A by a rank- ρ matrix.

Remark 5.1. For $\rho_+ = \rho$ we can write $T = T' = A^T G$ and skip Stage 2 because the matrix $A^T G$ is expected to serve as a desired approximate matrix basis by virtue of Theorem 5.2. In Appendix B we compute numerical rank by using randomized matrix multiplications instead of standard recipes that use orthogonalization or pivoting.

Remark 5.2. The increase of the dimension ρ_+ beyond the numerical rank ρ (called oversampling in [HMT11]) is relatively inexpensive if the bound ρ_+ is small. [HMT11] suggests using small oversampling even if the numerical rank ρ is known, because we have

$$\text{Probability } \{\|A - ATT^T\| \leq (1 + 9\sqrt{\rho_+ \min\{m, n\}})\sigma_{\rho+1}(A)\} \geq 1 - 3(\rho_+ - \rho)^{\rho - \rho_+} \text{ for } \rho_+ > \rho.$$

Theorem 5.2, however, bounds the norm $\|A - ATT^T\|$ strongly also for $\rho = \rho_+$, in good accordance with the data of Tables 7.2 and 7.3.

Remark 5.3. For a larger integer ρ we can substantially simplify Stage 1 of the algorithm by choosing structured multipliers G from the class of the subsample random Fourier transforms, called SRFTs. Under this choice the estimated probability of obtaining low rank approximation is close to the above case of Gaussian random multipliers G [T11]. Our tests in Section 7 provide informal empirical support for similar use of random Toeplitz multipliers G .

Remark 5.4. By applying rank revealing QR factorization at Stage 2 of the algorithm we can produce an orthogonal matrix T and consequently simplify Stage 3 by computing $\hat{A}_s = ATT^T$ (cf. (5.2)). We adopted such a variation of the algorithm in our tests in Section 7.

Remark 5.5. One can weaken reliability of the output to simplify Stage 3 by testing whether $\|K^T(A - \widehat{A}_\rho)L\| \leq \tau\|K\| \|A\| \|L\|$ for matrices $K \in \mathcal{G}_{0,1}^{m \times \rho'}$ and $L \in \mathcal{G}_{0,1}^{n \times \rho''}$ and for two small positive integers ρ' and ρ'' , possibly for $\rho' = \rho'' = 1$, instead of testing whether $\|\widehat{A}_\rho - A\| \leq \tau'\|(A)\|$. One can similarly simplify Stage 2.

Remark 5.6. Application of Proto-Algorithm 5.1 to the approximation of the leading singular spaces $\mathbb{T}_{\rho,A}$ and $\mathbb{S}_{\rho,A}$ is facilitated and its power is enhanced as the gaps increase between the singular values of the input matrix A . This motivates using the power transforms $A \implies B_h = (AA^T)^h A$ for positive integers h because $\sigma_j(B_h) = (\sigma_j(A))^{2h+1}$ for all j .

Remark 5.7. Clearly every matrix basis of the trailing singular space $\mathbb{T}_{A,n-\rho}$ of an $m \times n$ matrix A is orthogonal to every matrix basis of the leading singular space $\mathbb{T}_{\rho,A}$, and similarly for any pair of matrix bases of the spaces $\mathbb{S}_{A,m-\rho}$ and $\mathbb{S}_{\rho,a}$. One can exploit this duality for the computation and approximation of the bases.

6 Application to Tensor Train decomposition

Let

$$\mathbf{A} = [A(i_1, \dots, i_d)] \quad (6.1)$$

denote a d -dimensional tensor with entries $A(i_1, \dots, i_d)$ and spacial indices i_1, \dots, i_d ranging from 1 to n_1, \dots, n_d , respectively. Define the $d-1$ unfolding matrices $A_k = [A(i_1 \dots i_k; i_{k+1} \dots i_d)]$, $k = 1, \dots, d$, where the semicolon separates the multi-indices $i_1 \dots i_k$ and $i_{k+1} \dots i_d$, which define the rows and columns of the matrix A_k , respectively, $k = 1, \dots, d$. The paper [O09] proposed the following class of *Tensor Train Decompositions*, hereafter referred to as *TT Decompositions*, where the summation indices $\alpha_1, \dots, \alpha_{d-1}$ ranged from 1 to compression ranks r_1, \dots, r_{d-1} , respectively,

$$T = \sum_{\alpha_1, \dots, \alpha_{d-1}} G_1(i_1, \alpha_1)G_2(\alpha_1, i_1, \alpha_2) \cdots G_{d-1}(\alpha_{d-2}, i_{d-1}, \alpha_{d-1})G_d(\alpha_d, i_d). \quad (6.2)$$

Theorem 6.1. [O09]. For any tensor \mathbf{A} of (6.1) there exists a TT decomposition (6.2) such that $\mathbf{A} = \mathbf{T}$ and $r_k = \text{rank}(A_k)$ for $k = 1, \dots, d-1$.

There is a large and growing number of important applications of TT decompositions (6.2) to modern computations (cf. e.g., [OT09], [OT10], [OT11]) where the numerical ranks of the unfolding matrices A_k are much smaller than their ranks, and it is desired to compress TT decompositions respectively.

Theorem 6.2. [OT10]. For any tensor \mathbf{A} of (6.1) and any set of positive integers $r_k \leq \text{rank}(A_k)$, $k = 1, \dots, d-1$, there exists a TT decomposition (6.2) such that

$$\|\mathbf{A} - \mathbf{T}\|_F^2 \leq \sum_{k=1}^{d-1} \tau_k^2, \quad \tau_k = \min_{\text{rank}(B)=r_k} \|A_k - B\|_F, \quad k = 1, \dots, d-1. \quad (6.3)$$

The constructive proof of this theorem in [OT10] relies on inductive approximation of unfolding matrices by their SVDs truncated to the compression ranks r_k . Let us sketch this construction. For $d = 2$ we obtain a desired TT decomposition $T(i_1, i_2) = \sum_{\alpha_1}^{r_1} G_1(i_1, \alpha_1)G_2(\alpha_1, i_2)$ (that is a sum of r_1 outer products of r_1 pairs of vectors) simply by truncating the SVD of the matrix $A(i_1, i_2)$. At the inductive step one truncates the SVD of the first unfolding matrix $A_1 = S_{A_1} \Sigma_{A_1} T_{A_1}^T$ to obtain rank- r_1 approximation of this matrix $B_1 = S_{B_1} \Sigma_{B_1} T_{B_1}^T$ where $\Sigma_{B_1} = \text{diag}(\sigma_j(A_1))_{j=1}^{r_1}$ and the matrices S_{B_1} and T_{B_1} are formed by the first r_1 columns of the matrices S_{A_1} and T_{A_1} , respectively. Then it remains to approximate the tensor $\mathbf{B} = [B(i_1, \dots, i_d)]$ represented by the matrix B_1 . Rewrite it as $\sum_{\alpha_1=1}^{r_1} S_{B_1}(i_1; \alpha_1) \widehat{A}(\alpha_1; i_2 \dots i_d)$ for $\widehat{A} = \sum_{B_1} T_{B_1}^T$, represent \widehat{A} as the tensor $\widehat{\mathbf{A}} = [A(\alpha_1 i_2, i_3, \dots, i_d)]$ of dimension $d-1$, apply the inductive hypothesis to obtain a TT-approximation of this tensor, and extend it to a TT-approximation of the original tensor \mathbf{A} .

In [OT10] the authors specify this construction as their Algorithm 1, prove error norm bound (6.3), then point out that the “computation of the truncated SVD for large scale and possibly dense unfolding matrices ... is unaffordable in many dimensions”, propose “to replace SVD by some other dyadic decompositions $A_k \approx UV^T$, which can be computed with low complexity”, and finally specify such recipe as [OT10, Algorithm 2], which is an iterative algorithm for skeleton or pseudoskeleton decomposition of matrices and which they use at Stages 5 and 6 of their Algorithm 1. The cost of each iteration of [OT10, Algorithm 2] is quite low, and empirically the iteration converges fast, but the authors welcome alternative recipes having formal support.

Proto-Algorithm 5.1 can serve as an alternative to [OT10, Algorithm 2]. For the input matrix A_1 above we use $O(r_1)$ multiplications of this matrix by $O(r_1)$ vectors, which means a low computational cost for sparse and structured inputs, whereas the expected output is an approximate matrix basis for the space \mathbb{S}_{r_1, A_1} or \mathbb{T}_{r_1, A_1} and a rank- r_1 approximation to the matrix A_1 , within an expected error norm in $O(\sigma_{r_1+1}(A_1))$. This is the same order as in [OT10, Algorithm 1], but now we do not use SVDs. One can further decrease the error bound by means of small oversampling of Remark 5.2 and the power transform of Remark 5.6.

Remark 6.1. *A huge bibliography on tensor decompositions and on their application to fundamental matrix computations has been recently surveyed in [KB09], but with the omission of the early works [P72], [P79], [B80], [P84], [B86], where nontrivial tensor decompositions helped to accelerate the fundamental operation of matrix multiplication, probably the first application of this kind.*

7 Numerical Experiments

Our numerical experiments with random general, Hankel, Toeplitz and circulant matrices have been performed in the Graduate Center of the City University of New York on a Dell server with a dual core 1.86 GHz Xeon processor and 2G memory running Windows Server 2003 R2. The test Fortran code has been compiled with the GNU gfortran compiler within the Cygwin environment. Random numbers have been generated with the random_number intrinsic Fortran function, assuming the uniform probability distribution over the range $\{x : -1 \leq x < 1\}$. The tests have been designed by the first author and performed by his coauthor.

7.1 GENP with random circulant multipliers

Table 7.1 shows the results of our tests of the solution of a nonsingular well conditioned linear system $A\mathbf{y} = \mathbf{b}$ of n equations whose coefficient matrix has ill conditioned $n/2 \times n/2$ leading principal block for $n = 64, 256, 1024$. We have performed 100 numerical tests for each dimension n and computed the maximum, minimum and average relative residual norms $\|A\mathbf{y} - \mathbf{b}\|/\|\mathbf{b}\|$ as well as standard deviation. GENP applied to these systems outputs corrupted solutions with residual norms ranging from 10 to 10^8 . When we preprocessed the systems with circulant multipliers filled with ± 1 (choosing the n signs \pm at random), the norms decreased to at worst 10^{-7} for all inputs. Table 7.1 also shows further decrease of the norm in a single step of iterative refinement. Table 2 in [PQZa] shows similar results of the tests where the input matrices have been chosen similarly but so that their every leading $k \times k$ block had numerical rank k or $k - 1$ and where Householder multipliers $I_n - \mathbf{u}\mathbf{v}^T/\mathbf{u}^T\mathbf{v}$ replaced the circulant multipliers. Here \mathbf{u} and \mathbf{v} denote two vectors filled with integers 1 and -1 under random choice of the signs $+$ and $-$.

7.2 Approximation of the tails and heads of SVDs and low-rank approximation of a matrix

At some specified stages of our tests of this subsection we performed additions, subtractions and multiplications with infinite precision (hereafter referred to as *error-free ring operations*). At the other stages we performed computations with double precision, and we rounded to double precision all random values. We performed at most two refinement iterations for the computed solution of every linear system of equations and matrix inverse.

Tables 7.2 and 7.3 display the data from our tests on the approximation of leading singular spaces of the SVD of an $n \times n$ matrix A having numerical rank q and on the approximation of this matrix with a matrix of rank ρ . For $n = 64, 128, 256$ and $\rho = 1, 8, 32$ we generated $n \times n$ random orthogonal matrices S and T and diagonal matrices $\Sigma = \text{diag}(\sigma_j)_{j=1}^n$ such that $\sigma_j = 1/j$, $j = 1, \dots, \rho$, $\sigma_j = 10^{-10}$, $j = \rho + 1, \dots, n$ (cf. [H02, Section 28.3]). Then we applied error-free ring operations to compute the input matrices $A = S_A \Sigma_A T_A^T$, for which $\|A\| = 1$ and $\kappa(A) = 10^{10}$. Furthermore we generated random $n \times \rho$ matrices G (for $\rho = 1, 8, 32$) and successively computed the matrices $B_{\rho,A} = A^T G$, $T_{\rho,A}$, $B_{\rho,A} Y_{\rho,A}$ as a least-squares approximation to $T_{\rho,A}$, $Q_{\rho,A} = Q(B_{\rho,A})$ (cf. Fact 2.2), and $A - A Q_{\rho,A} (Q_{\rho,A})^T$ (by applying error-free ring operations). Table 7.2 summarizes the data on the residual norms $\text{rn}^{(1)} = \|B_{\rho,A} Y_{\rho,A} - T_{\rho,A}\|$ and $\text{rn}^{(2)} = \|A - A Q_{\rho,A} (Q_{\rho,A})^T\|$ obtained in 100 runs of our tests for every pair of n and ρ .

We have also performed similar tests where we generated random Toeplitz $n \times \rho$ matrices T (for $\rho = 8, 32$) and then replaced the above approximate matrix bases $B_{\rho,A} = A^T G$ for the leading singular space $\mathbb{T}_{\rho,A}$ by the matrices $B_{\rho,A} = A^T T$. Table 7.3 displays the results of these tests. In both Tables 7.2 and 7.3 the residual norms are more or less equally small.

Table 7.1: Relative residual norms: randomized circulant GENP for well conditioned linear systems with ill conditioned leading blocks (cf. [PQZa, Table 2])

dimension	iterations	min	max	mean	std
64	0	4.7×10^{-14}	8.0×10^{-11}	4.0×10^{-12}	1.1×10^{-11}
64	1	1.9×10^{-15}	5.3×10^{-13}	2.3×10^{-14}	5.4×10^{-14}
256	0	1.7×10^{-12}	1.4×10^{-7}	2.0×10^{-9}	1.5×10^{-8}
256	1	8.3×10^{-15}	4.3×10^{-10}	4.5×10^{-12}	4.3×10^{-11}
1024	0	1.7×10^{-10}	4.4×10^{-9}	1.4×10^{-9}	2.1×10^{-9}
1024	1	3.4×10^{-14}	9.9×10^{-14}	6.8×10^{-14}	2.7×10^{-14}

Table 7.2: Heads of SVDs and low-rank approximation by using random multipliers G

q	rn_i	n	min	max	mean	std
1	$\text{rn}^{(1)}$	64	2.35×10^{-10}	1.32×10^{-07}	3.58×10^{-09}	1.37×10^{-08}
1	$\text{rn}^{(1)}$	128	4.41×10^{-10}	3.28×10^{-08}	3.55×10^{-09}	5.71×10^{-09}
1	$\text{rn}^{(1)}$	256	6.98×10^{-10}	5.57×10^{-08}	5.47×10^{-09}	8.63×10^{-09}
1	$\text{rn}^{(2)}$	64	8.28×10^{-10}	1.32×10^{-07}	3.86×10^{-09}	1.36×10^{-08}
1	$\text{rn}^{(2)}$	128	1.21×10^{-09}	3.28×10^{-08}	3.91×10^{-09}	5.57×10^{-09}
1	$\text{rn}^{(2)}$	256	1.74×10^{-09}	5.58×10^{-08}	5.96×10^{-09}	8.47×10^{-09}
8	$\text{rn}^{(1)}$	128	2.56×10^{-09}	1.16×10^{-06}	4.30×10^{-08}	1.45×10^{-07}
8	$\text{rn}^{(1)}$	256	4.45×10^{-09}	3.32×10^{-07}	3.40×10^{-08}	5.11×10^{-08}
8	$\text{rn}^{(2)}$	64	1.46×10^{-09}	9.56×10^{-08}	5.77×10^{-09}	1.06×10^{-08}
8	$\text{rn}^{(2)}$	128	1.64×10^{-09}	4.32×10^{-07}	1.86×10^{-08}	5.97×10^{-08}
8	$\text{rn}^{(2)}$	256	2.50×10^{-09}	1.56×10^{-07}	1.59×10^{-08}	2.47×10^{-08}
32	$\text{rn}^{(1)}$	64	6.80×10^{-09}	2.83×10^{-06}	1.01×10^{-07}	3.73×10^{-07}
32	$\text{rn}^{(1)}$	128	1.25×10^{-08}	6.77×10^{-06}	1.28×10^{-07}	6.76×10^{-07}
32	$\text{rn}^{(1)}$	256	1.85×10^{-08}	1.12×10^{-06}	1.02×10^{-07}	1.54×10^{-07}
32	$\text{rn}^{(2)}$	64	1.84×10^{-09}	6.50×10^{-07}	2.30×10^{-08}	8.28×10^{-08}
32	$\text{rn}^{(2)}$	128	3.11×10^{-09}	1.45×10^{-06}	2.87×10^{-08}	1.45×10^{-07}
32	$\text{rn}^{(2)}$	256	4.39×10^{-09}	2.16×10^{-07}	2.37×10^{-08}	3.34×10^{-08}

Table 7.3: Heads of SVDs and low-rank approximations by using random Toeplitz multipliers T

q	$\text{rrn}^{(i)}$	n	min	max	mean	std
8	$\text{rrn}^{(1)}$	64	2.22×10^{-09}	7.89×10^{-06}	1.43×10^{-07}	9.17×10^{-07}
8	$\text{rrn}^{(1)}$	128	3.79×10^{-09}	4.39×10^{-05}	4.87×10^{-07}	4.39×10^{-06}
8	$\text{rrn}^{(1)}$	256	5.33×10^{-09}	3.06×10^{-06}	6.65×10^{-08}	3.12×10^{-07}
8	$\text{rrn}^{(2)}$	64	1.13×10^{-09}	3.66×10^{-06}	6.37×10^{-08}	4.11×10^{-07}
8	$\text{rrn}^{(2)}$	128	1.81×10^{-09}	1.67×10^{-05}	1.90×10^{-07}	1.67×10^{-06}
8	$\text{rrn}^{(2)}$	256	2.96×10^{-09}	1.25×10^{-06}	2.92×10^{-08}	1.28×10^{-07}
32	$\text{rrn}^{(1)}$	64	6.22×10^{-09}	5.00×10^{-07}	4.06×10^{-08}	6.04×10^{-08}
32	$\text{rrn}^{(1)}$	128	2.73×10^{-08}	4.88×10^{-06}	2.57×10^{-07}	8.16×10^{-07}
32	$\text{rrn}^{(1)}$	256	1.78×10^{-08}	1.25×10^{-06}	1.18×10^{-07}	2.03×10^{-07}
32	$\text{rrn}^{(2)}$	64	1.64×10^{-09}	1.26×10^{-07}	9.66×10^{-09}	1.48×10^{-08}
32	$\text{rrn}^{(2)}$	128	5.71×10^{-09}	9.90×10^{-07}	5.50×10^{-08}	1.68×10^{-07}
32	$\text{rrn}^{(2)}$	256	4.02×10^{-09}	2.85×10^{-07}	2.74×10^{-08}	4.48×10^{-08}

8 Conclusions

It is well known that random matrices tend to be well conditioned, and this property motivates our application of random matrix multipliers for advancing some fundamental matrix computations. We first prove the basic fact that with a probability close to 1 multiplication by a Gaussian random matrix does not increase the condition number of a matrix and of its any block dramatically compared to the condition number of the input matrix. As an immediate implication random multipliers are likely to stabilize numerically GENP (that is Gaussian elimination with no pivoting) and block Gaussian elimination applied to a nonsingular and well conditioned matrix, possibly having ill conditioned and singular leading blocks. by applying to input matrix randomized structured multipliers. Another basic fact states that with a probability close to 1 the column sets of the products $A^T G$ and AH where an $m \times n$ matrix A has a numerical rank ρ and G and H are Gaussian random matrices of sizes $m \times \rho$ and $n \times \rho$, respectively, approximate some bases for the left and right leading singular spaces $\mathbb{S}_{\rho,A}$ and $\mathbb{T}_{\rho,A}$ associated with the ρ largest singular values of the matrix A . Having any of such approximate bases available we can readily approximate the matrix A by a matrix of rank ρ . This has further well known extensions to many important matrix computations, and we point out new extensions to the approximation of a matrix by a structured matrix lying nearby, to computing numerical rank of a matrix, and to Tensor Train approximation.

Our tests consistently showed efficiency of the proposed techniques even where instead of general Gaussian random multipliers we applied structured and sparse multipliers. In these cases randomization was limited to much fewer random parameters or just to the choice of the signs \pm of a few auxiliary vectors. The recent paper [T11] is an important step toward understanding and exploiting this phenomenon and should motivate further research effort. Another natural research subject is the combination of randomized matrix multiplication with randomized techniques of additive preprocessing and augmentation, recently studied in [PGMQ], [PIMR10], [PQ10], [PQ12], [PQZC], [PQZa], [PQZb], and [PQZc].

Appendix

A Uniform random sampling and nonsingularity of random matrices

Uniform random sampling of elements from a finite set Δ is their selection from this set at random, independently of each other and under the uniform probability distribution on the set Δ .

Theorem A.1. *Under the assumptions of Lemma 3.1 let the values of the variables of the polynomial be randomly and uniformly sampled from a finite set Δ . Then the polynomial vanishes with a probability at most $\frac{d}{|\Delta|}$.*

Corollary A.1. *Let the entries of a general or Toeplitz $m \times n$ matrix have been randomly and uniformly sampled from a finite set Δ of cardinality $|\Delta|$ (in any fixed ring). Let $l = \min\{m, n\}$. Then (a) every $k \times k$ submatrix M for $k \leq l$ is nonsingular with a probability at least $1 - \frac{k}{|\Delta|}$ and (b) is strongly nonsingular with a probability at least $1 - \sum_{i=1}^k \frac{i}{|\Delta|} = 1 - \frac{(k+1)k}{2|\Delta|}$.*

Proof. The claimed properties of nonsingularity and nonvanishing hold for generic matrices. The singularity of a $k \times k$ matrix means that its determinant vanishes, but the determinant is a polynomial of total degree k in the entries. Therefore Theorem A.1 implies parts (a) and consequently (b). Part (c) follows because a fixed entry of the inverse vanishes if and only if the respective entry of the adjoint vanishes, but up to the sign the latter entry is the determinant of a $(k-1) \times (k-1)$ submatrix of the input matrix M , and so it is a polynomial of degree $k-1$ in its entries. \square

B Computation of numerical ranks

The customary algorithms for the numerical rank of a matrix rely on computing its SVD or rank revealing factorization and involving pivoting or orthogonalization. Proto-Algorithm 5.1 also uses rank revealing factorization at Stage 2 and matrix inversion or orthogonalization at Stage 3, but only with matrices of small sizes provided the integer ρ_+ is small. Our next alternative algorithm avoids pivoting and orthogonalization even where the numerical rank ρ is large. As by-product we compute an approximate matrix basis within an error norm in $O(\sigma_{\rho+1}(\tilde{A}))$ for the leading singular space $\mathbb{T}_{\rho, \tilde{A}}$ of an $m \times n$ matrix \tilde{A} and can extend this readily to computing a rank- ρ approximation of the matrix \tilde{A} (see Remark B.1). We let $m \geq n$ (else shift to \tilde{A}^T), let $[\rho_-, \rho_+] = [0, n]$ unless we know a more narrow range, and successively test the selected candidate integers in the range $[\rho_-, \rho_+]$ until we find the numerical rank ρ . To improve reliability, we can repeat the tests for distinct values of random parameters (see Remarks B.1).

Exhaustive search defines and verifies the numerical rank ρ with probability near 1, but with proper policies one can use fewer and simpler tests because for $G \in \mathcal{G}_{0,1}^{m \times s}$ (and empirically for various random sparse and structured matrices G as well) the matrix $B = \tilde{A}^T G$ is expected (a) to have full rank and to be well conditioned if and only if $s \geq \rho$, (b) to approximate a matrix basis (within an error norm in $O(\sigma_{\rho+1}(\tilde{A}))$) for a linear space $\mathbb{T} \supseteq \mathbb{T}_{\rho, B} = \mathbb{T}_{\rho, \tilde{A}}$ where $s \geq \rho$, and (c) to approximate a matrix basis (within an error norm in $O(\sigma_{\rho+1}(\tilde{A}))$) for the space $\mathbb{T}_{\rho, \tilde{A}}$ where $s = \rho$. Property (a) is implied by Theorem 4.1, properties (b) and (c) by Theorem 5.2.

Proto-Algorithm B.1. Numerical rank without pivoting and orthogonalization (see Remarks B.1–B.4).

INPUT: *Two integers ρ_- and ρ_+ and a matrix $\tilde{A} \in \mathbb{R}^{m \times n}$ having unknown numerical rank $\rho = \text{rank}(\tilde{A})$ in the range $[\rho_-, \rho_+]$ such that $0 \leq \rho_- < \rho_+ \leq n \leq m$, a rule for the selection of a candidate integer ρ in a range $[\rho_-, \rho_+]$, and a Subroutine COND that determines whether a given matrix has full rank and is well conditioned or not.*

OUTPUT: *an integer ρ expected to equal numerical rank of the matrix \tilde{A} and a matrix B expected to approximate (within an error norm in $O(\sigma_{\rho+1}(\tilde{A}))$) a matrix basis of the singular space $\mathbb{T}_{\rho, \tilde{A}}$. (Both expectations can actually fail, but with a low probability, see Remark B.1.)*

INITIALIZATION: *Generate matrix $G \in \mathcal{G}_{0,1}^{m \times \rho_+}$ and write $B = \tilde{A}$, $G_\rho = G(I_\rho \mid O_{\rho, m-\rho})^T$ for $\rho = \rho_-, \rho_- + 1, \dots, \rho_+$ (The $m \times \rho$ matrix G_ρ is formed by the first ρ columns of the matrix G .)*

COMPUTATIONS:

1. *Stop and output $\rho = \rho_+$ and the matrix B if $\rho_- = \rho_+$. Otherwise fix an integer ρ in the range $[\rho_-, \rho_+]$.*

2. Compute the matrix $B' = B^T G_\rho$ and apply to it the Subroutine COND.
3. If this matrix has full rank and is well conditioned, write $\rho_+ = \rho$ and $B = B'$ and go to Stage 1. Otherwise write $\rho_- = \rho$ and go to Stage 1.

Remark B.1. The algorithm can output a wrong value of the numerical rank, although by virtue of Theorems 5.1 and 5.2 combined this occurs with a low probability. One can decrease this probability by reapplying the algorithm to the same inputs and choosing distinct random parameters. Furthermore one can fix a tolerance τ of order $\sigma_{\rho+1}(\tilde{A})$, set $T = B$, and apply Stage 3 of Proto-Algorithm 5.1. Then $\text{nrank}(\tilde{A})$ is expected to exceed the computed value ρ if this stage outputs FAILURE and to equal ρ otherwise, in which case Proto-Algorithm 5.1 also outputs a rank- ρ approximation of the matrix \tilde{A} (within an error norm $\tau\|\tilde{A}\|$ in $O(\sigma_{\rho+1}(\tilde{A}))$). For a sufficiently small tolerance τ the latter outcome implies that certainly $\rho \geq \text{nrank}(\tilde{A})$.

Remark B.2. We can avoid pivoting and orthogonalization in Subroutine COND by applying the Power or Lanczos algorithms [B74], [D83], [KW92]. We can first apply the Power Method to the matrix $S = \tilde{A}^T \tilde{A}$ or $S = \tilde{A} \tilde{A}^T$ to yield a close upper bound σ_+^2 on its largest eigenvalue $\sigma_1^2(\tilde{A})$ and then to the matrix $\sigma_+^2 I - \tilde{A}^T \tilde{A}$ to approximate the smallest eigenvalue of the matrix S , equal to $\sigma_n^2(\tilde{A})$. The Lanczos algorithm approximates both extremal eigenvalues of the matrix $S = \tilde{A}^T \tilde{A}$ (equal to $\sigma_1^2(\tilde{A})$ and $\sigma_l^2(\tilde{A})$, $l = \min\{m, n\}$ respectively) and converges much faster [GL96, Sections 9.1.4, 9.1.5].

Remark B.3. One can simplify Stage 2 by applying the Subroutine COND to the matrix $G'_\rho = F_\rho G_\rho$ of a smaller size (rather than to G_ρ) where $F_\rho \in \mathcal{G}_{0,1}^{\rho \times m}$. By virtue of Theorem 4.1 the matrices G_ρ and G'_ρ are likely to have condition numbers of the same order.

Remark B.4. The binary search $\rho = \lceil (\rho_- + \rho_+) / 2 \rceil$ is an attractive policy for choosing the candidate values ρ , but one may prefer to move toward ρ_- , the left end of the range more rapidly, to decrease the size of the matrix B' .

Acknowledgements: Our research has been supported by NSF Grant CCF-1116736 and PSC CUNY Awards 64512-0042 and 65792-0043. We are also grateful to Mr. Jesse Wolf for helpful comments.

References

- [A94] O. Axelsson, *Iterative Solution Methods*, Cambridge Univ. Press, England, 1994.
- [B74] D.W. Boyd, The Power method for l^p Norms, *Linear Algebra and Its Applications*, **9**, 95–101, 1974.
- [B80] D. Bini, Border Rank of $p \times q \times 2$ Tensors and the Optimal Approximation of a Pair of Bilinear Forms, in *Lecture Notes in Computer Science*, **85**, 98–108, Springer, 1980.
- [B86] D. Bini, Border Rank of $m \times n \times (mn - q)$ Tensors, *Linear Algebra and Its Applications*, **79**, 45–51, 1986.
- [B02] M. Benzi, Preconditioning Techniques for Large Linear Systems: a Survey, *J. of Computational Physics*, **182**, 418–477, 2002.
- [B11] C. Beltrán, Estimates on the Condition Number of Random, Rank-deficient Matrices, *IMA Journal of Numerical Analysis*, **31**, 1, 25–39, 2011.
- [BM01] D. A. Bini, B. Meini, Approximate Displacement Rank and Applications, in *AMS Conference "Structured Matrices in Operator Theory, Control, Signal and Image Processing"*, Boulder, 1999 (edited by V. Olshevsky), *American Math. Society*, 215–232, Providence, RI, 2001.

- [CD05] Z. Chen, J. J. Dongarra, Condition Numbers of Gaussian Random Matrices, *SIAM J. on Matrix Analysis and Applications*, **27**, 603–620, 2005.
- [D83] J. D. Dixon, Estimating Extremal Eigenvalues and Condition Numbers of Matrices, *SIAM J. on Numerical Analysis*, **20**, 4, 812–814, 1983.
- [D88] J. Demmel, The Probability That a Numerical Analysis Problem Is Difficult, *Math. of Computation*, **50**, 449–480, 1988.
- [DL78] R. A. Demillo, R. J. Lipton, A Probabilistic Remark on Algebraic Program Testing, *Information Processing Letters*, **7**, 4, 193–195, 1978.
- [DS01] K. R. Davidson, S. J. Szarek, Local Operator Theory, Random Matrices, and Banach Spaces, in *Handbook on the Geometry of Banach Spaces* (W. B. Johnson and J. Lindenstrauss editors), pages 317–368, North Holland, Amsterdam, 2001.
- [E88] A. Edelman, Eigenvalues and Condition Numbers of Random Matrices, *SIAM J. on Matrix Analysis and Applications*, **9**, 4, 543–560, 1988.
- [ES05] A. Edelman, B. D. Sutton, Tails of Condition Number Distributions, *SIAM J. on Matrix Analysis and Applications*, **27**, 2, 547–560, 2005.
- [G97] A. Greenbaum, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, 1997.
- [GE96] M. Gu, S. C. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM Journal on Scientific Computing*, **17**, 848–869, 1996.
- [GL96] G. H. Golub, C. F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, Maryland, 1996 (third addition).
- [H02] N. J. Higham, *Accuracy and Stability in Numerical Analysis*, SIAM, Philadelphia, 2002 (second edition).
- [HMT11] N. Halko, P. G. Martinsson, J. A. Tropp, Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions, *SIAM Review*, **53**, 2, 217–288, 2011.
- [HP92] Y. P. Hong, C.-T. Pan, The rank revealing QR decomposition and SVD, *Math. of Computation*, **58**, 213–232, 1992.
- [KB09] T. G. Kolda, B. W. Bader, Tensor Decompositions and Applications, *SIAM Review*, **51**, 3, 455–500, 2009.
- [KKM79] T. Kailath, S. Y. Kung, M. Morf, Displacement Ranks of Matrices and Linear Equations, *Journal Math. Analysis and Appls*, **68**, 2, 395–407, 1979.
- [KL94] C. S. Kenney, A. J. Laub, Small-Sample Statistical Condition Estimates for General Matrix Functions, *SIAM J. on Scientific and Statistical Computing*, **15**, 36–61, 1994.
- [KW92] J. Kuczynski, H. Wozniakowski, Estimating the Largest Eigenvalue by the Power and Lanczos Algorithms with a Random Start, *SIAM J. on Matrix Analysis and Applications*, **13**, 1094–1122, 1992.
- [O09] I. V. Oseledets, A Compact Matrix Form of the d-Dimensional Tensor Decomposition, Preprint 2009-01, *INM RAS*, March 2009.
- [OT09] I. V. Oseledets, E. E. Tyrtyshnikov, Breaking the Curse of Dimensionality, or How to Use SVD in Many Dimensions, *SIAM J. Scientific Comp.*, **31**, 5, 3744–3759, 2009.

- [OT10] I. V. Oseledets, E. E. Tyrtyshnikov, TT-cross Approximation for Multidimensional Arrays, *Linear Algebra Appls.* **432**, **1**, 70–88, 2010.
- [OT11] I. Oseledets, E. E. Tyrtyshnikov, Algebraic Wavelet Transform via Quantics Tensor Train Decomposition, *SIAM J. Sci. Comp.*, **33**, **3**, 1315–1328, 2011.
- [P72] V. Y. Pan, On Schemes for the Evaluation of Products and Inverses of Matrices (in Russian), *Uspekhi Matematicheskikh Nauk*, **27**, **5 (167)**, 249–250, 1972.
- [P79] V. Y. Pan, Fields Extension and Trilinear Aggregating, Uniting and Canceling for the Acceleration of Matrix Multiplication, *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science (FOCS 79)*, 28–38, IEEE Computer Society Press, Long Beach, California, 1979.
- [P84] V. Y. Pan, How Can We Speed up Matrix Multiplication? *SIAM Review*, **26**, **3**, 393–415, 1984.
- [P00a] C.-T. Pan, On the Existence and Computation of Rank-revealing LU Factorization, *Linear Algebra and Its Applications*, **316**, 199–222, 2000.
- [P01] V. Y. Pan, *Structured Matrices and Polynomials: Unified Superfast Algorithms*, Birkhäuser/Springer, Boston/New York, 2001.
- [P10] V. Y. Pan, Newton’s Iteration for Matrix Inversion, Advances and Extensions, pp. 364–381, in *Matrix Methods: Theory, Algorithms and Applications* (dedicated to the Memory of Gene Golub, edited by V. Olshevsky and E. Tyrtyshnikov), World Scientific Publishing, New Jersey, ISBN-13 978-981-283-601-4, ISBN-10-981-283-601-2 (2010).
- [PGMQ] V. Y. Pan, D. Grady, B. Murphy, G. Qian, R. E. Rosholt, A. Ruslanov, Schur Aggregation for Linear Systems and Determinants, *Theoretical Computer Science, Special Issue on Symbolic–Numerical Algorithms* (D. A. Bini, V. Y. Pan, and J. Verschelde editors), **409**, **2**, 255–268, 2008.
- [PIMR10] V. Y. Pan, D. Ivolgin, B. Murphy, R. E. Rosholt, Y. Tang, X. Yan, Additive Preconditioning for Matrix Computations, *Linear Algebra and Its Applications*, **432**, 1070–1089, 2010.
- [PQ10] V. Y. Pan, G. Qian, Randomized Preprocessing of Homogeneous Linear Systems of Equations, *Linear Algebra and Its Applications*, **432**, 3272–3318, 2010.
- [PQ12] V. Y. Pan, G. Qian, Solving Linear Systems of Equations with Randomization, Augmentation and Aggregation, *Linear Algebra and Its Applications*, **437**, 2851–1876, 2012.
- [PQa] V. Y. Pan, G. Qian, Condition Numbers of Random Toeplitz and Circulant Matrices, Tech. Report TR 2012013, *PhD Program in Comp. Sci., Graduate Center, CUNY*, 2012.
Available at <http://www.cs.gc.cuny.edu/tr/techreport.php?id=442>
- [PQZa] V. Y. Pan, G. Qian, A. Zheng, Randomized Preprocessing versus Pivoting, *Linear Algebra and Its Applications*, in print. <http://dx.doi.org/10.1016/j.laa.2011.02.052>
- [PQZb] V. Y. Pan, G. Qian, A. Zheng, Randomized Matrix Computations, Tech. Report TR 2012009, *PhD Program in Comp. Sci., Graduate Center, CUNY*
Available at <http://www.cs.gc.cuny.edu/tr/techreport.php?id=438> and <http://arxiv.org/abs/1210.7476>

- [PQZc] V. Y. Pan, G. Qian, A. Zheng, Randomized Augmentation and Additive Preprocessing, Tech. Report TR 201201x, *PhD Program in Comp. Sci., Graduate Center, CUNY* 2012.
- [PQZC] V. Y. Pan, G. Qian, A. Zheng, Z. Chen, Matrix Computations and Polynomial Root-finding with Preprocessing, *Linear Algebra and Its Applications*, **434**, 854–879, 2011.
- [PY09] V. Y. Pan, X. Yan, Additive Preconditioning, Eigenspaces, and the Inverse Iteration, *Linear Algebra and Its Applications*, **430**, 186–203, 2009.
- [S80] J. T. Schwartz, Fast Probabilistic Algorithms for Verification of Polynomial Identities, *Journal of ACM*, **27**, **4**, 701–717, 1980.
- [S98] G. W. Stewart, *Matrix Algorithms, Vol I: Basic Decompositions*, SIAM, 1998.
- [SST06] A. Sankar, D. Spielman, S.-H. Teng, Smoothed Analysis of the Condition Numbers and Growth Factors of Matrices, *SIAM J. on Matrix Analysis*, **28**, **2**, 446–476, 2006.
- [T11] J. A. Tropp, Improved analysis of the subsampled randomized Hadamard transform, *Adv. Adapt. Data Anal.*, **3**, **1–2**, Special Issue, "Sparse Representation of Data and Images," 115-126, 2011.
- [Z79] R. E. Zippel, Probabilistic Algorithms for Sparse Polynomials, *Proceedings of EU-ROSAM'79, Lecture Notes in Computer Science*, **72**, 216–226, Springer, Berlin, 1979.