2013

TR-2013016: Supporting GENP with Random Multipliers

Victor Y. Pan

Guoliang Qian

Xiaodong Yan

Supporting GENP with Random Multipliers

Victor Y. Pan[1,2], Guoliang Qian[2], and Xiaodong Yan[2]

[1] Department of Mathematics and Computer Science
Lehman College of the City University of New York
Bronx, NY 10468 USA
victor.pan@lehman.cuny.edu,
WWW home page: http://comet.lehman.cuny.edu/vpan/
[2] Ph.D. Programs in Mathematics and Computer Science
The Graduate Center of the City University of New York
New York, NY 10036 USA
gqian@gc.cuny.edu, xyannyc@yahoo.com

Abstract. We prove that standard Gaussian random multipliers are expected to stabilize numerically both Gaussian elimination with no pivoting and block Gaussian elimination. Our tests show similar results where we applied circulant random multipliers instead of Gaussian ones.

Key Words: Random matrices, Random multipliers, GENP

1 Introduction

It is well known that even for a nonsingular and well conditioned input matrix, Gaussian elimination fails in numerical computations with rounding errors as soon as it encounters a vanishing or nearly vanishing leading (that is northwestern) entry. In practice the users avoid such encounters by applying GEPP, which stands for Gaussian elimination with partial pivoting and has some limited formal but ample empirical support. Partial pivoting, that is an appropriate row interchange, however, takes its toll: "pivoting usually degrades the performance" [GL96, page 119]. It interrupts the stream of arithmetic operations with foreign operations of comparison, involves book-keeping, compromises data locality, increases communication overhead and data dependence, and tends to destroy matrix structure. The users apply *Gaussian elimination with no pivoting* (hereafter we refer to it as *GENP*) or its variants wherever this application is numerically safe, in particular where a well conditioned input matrix is positive definite, diagonally dominant, or totally positive. By following the idea in [PGMQ, Section 12.2], briefly revisited in [PQZ13, Section 6], we generalize this class to all nonsingular and well conditioned matrices by *preconditioning* GENP with standard Gaussian random multipliers (hereafter we call them just *Gaussian*). We provide formal support for this approach, which can be extended to block Gaussian elimination (see Section 2.4). Our tests are in good accordance with our formal study. Moreover multiplication by circulant (rather than Gaussian) random multipliers has the same empirical power, although it uses by a

factor of n fewer random parameters and by a factor of $n/\log(n)$ fewer flops in the case of $n \times n$ inputs. The latter saving is particularly dramatic in the important case where the input matrix has Toeplitz structure. Our study is a new demonstration of the power of randomized matrix algorithms (cf. [M11], [HMT11], and the bibliography therein).

2 Some definitions and basic results

We assume computations in the field $\mathbb{R}$ of real numbers, but the extension to the case of the complex field $\mathbb{C}$ is quite straightforward. Hereafter "flop" stands for "arithmetic operation", and "Gaussian matrix" stands for "standard Gaussian random matrix". The concepts "large", "small", "near", "closely approximate", "ill conditioned" and "well conditioned" are quantified in the context. By saying "expect" and "likely" we mean "with probability 1 or close to 1" (we do not use the concept of the expected value). Next we recall and extend some customary definitions of matrix computations [GL96], [S98].

2.1 Some basic definitions of matrix computations

$\mathbb{R}^{m\times n}$ is the class of real $m \times n$ matrices $A = (a_{i,j})_{i,j}^{m,n}$. $(B_1 \mid \ldots \mid B_k)$ is a $1 \times k$ block matrix with the blocks $B_1, \ldots, B_k$. $\text{diag}(B_1, \ldots, B_k) = \text{diag}(B_j)_{j=1}^{k}$ is a $k \times k$ block diagonal matrix with the diagonal blocks $B_1, \ldots, B_k$. In both cases the blocks B_j can be rectangular. $\mathbf{e}_i$ is the ith coordinate vector of dimension n for $i = 1, \ldots, n$. These vectors define the $n \times n$ identity matrix $I_n = (\mathbf{e}_1 \mid \ldots \mid \mathbf{e}_n)$. $O_{k,l}$ is the $k \times l$ matrix filled with zeros. We write I and O where the matrix size is defined by context. $\text{rank}(A)$ denotes the rank of a matrix A. A^T is its transpose. $A_{k,l}$ is its leading, that is northwestern $k \times l$ block submatrix, and in Section 2.4 we also write $A^{(k)} = A_{k,k}$. A_s^T denotes the transpose $(A_s)^T$ of a matrix A_s, e.g., $A_{k,l}^T$ stands for $(A_{k,l})^T$. A matrix of a rank ρ has *generic rank profile* if all its leading $i \times i$ blocks are nonsingular for $i = 1, \ldots, \rho$. If such a matrix is nonsingular itself, then it is called *strongly nonsingular*. *Preprocessing* $A \to FA$, $A \to AH$, and $A \to FAH$, for nonsingular matrices F and H, reduces the inversion of a matrix A to the inversion of the products FA, AH, or FAH, and similarly for the solution of a linear system of equations.

Fact 1. *Assume three nonsingular matrices F, A, and H and a vector $\mathbf{b}$. Then $A^{-1} = H(FAH)^{-1}F$, $FAH\mathbf{y} = F\mathbf{b}$, $\mathbf{x} = H\mathbf{y}$ if $A\mathbf{x} = \mathbf{b}$.*

2.2 Matrix norms, orthogonality, SVD, and pseudo-inverse

$||A|| = ||A||_2 = \sup_{\mathbf{v}^T\mathbf{v}=1} ||A\mathbf{v}||$ is the spectral norm of a matrix $A = (a_{i,j})_{i,j=1}^{m,n}$, $||AB|| \le ||A||\,||B||$. A real matrix Q is called *orthogonal* if $Q^TQ = I$ or $QQ^T = I$.

An *SVD* or *full SVD* of an $m \times n$ matrix A of a rank ρ is a factorization

$$A = S_A \Sigma_A T_A^T \tag{1}$$

where $S_A = (\mathbf{s}_i)_{i=1}^m$ and $T_A = (\mathbf{t}_j)_{j=1}^n$ are square orthogonal matrices, $\Sigma_A = \mathrm{diag}(\widehat{\Sigma}_A, O_{m-\rho,n-\rho})$, $\widehat{\Sigma}_A = \mathrm{diag}(\sigma_j(A))_{j=1}^\rho$, $\sigma_j = \sigma_j(A) = \sigma_j(A^T)$ is the jth largest singular value of a matrix A, and $\sigma_j = 0$ for $j > \rho$, $\sigma_\rho > 0$, $\sigma_1 = \max_{||\mathbf{x}||=1} ||A\mathbf{x}|| = ||A||$, and

$$\min_{\mathrm{rank}(B)\le s-1} \quad ||A - B|| = \sigma_s(A),\ s = 1, 2, \ldots, \tag{2}$$

Fact 2. *If A_0 is a submatrix of a matrix A, then $\sigma_j(A) \ge \sigma_j(A_0)$ for all j.*

Proof. Extend [GL96, Corollary 8.6.3].

Fact 3. *(Cf. [GL96, Corollary 8.6.3].) Suppose $r + l \le n \le m$, $l \ge 0$, $1 \le k \le r$, $A \in \mathbb{R}^{m\times n}$, and $A_{m,r}$ is the leftmost $m \times r$ block of the matrix A. Then $\sigma_k(A_{m,r}) \ge \sigma_{k+l}(A_{m,r+l})$.*

$A^+ = T_A\,\mathrm{diag}(\widehat{\Sigma}_A^{-1}, O_{n-\rho,m-\rho})S_A^T$ is the Moore–Penrose pseudo-inverse of the matrix A of (1). A^{+T} stands for $(A^+)^T = (A^T)^+$, and A_s^+ stands for $(A_s)^+$, e.g., $A_{k,l}^+$ denotes $(A_{k,l})^+$.

If a matrix A has full column rank ρ, then $A^+ = (A^TA)^{-1}A^T$ and

$$||A^+|| = 1/\sigma_\rho(A). \tag{3}$$

Corollary 1. *Assume that* $\mathrm{rank}(A_{m,r}) = r$ *and* $\mathrm{rank}(A_{m,r+l}) = r + l$ *for the matrices $A_{m,r}$ and $A_{m,r+l}$ of Fact 3. Then $||A_{m,r}^+|| \le ||A_{m,r+l}^+||$.*

Proof. Combine Fact 3 for $k = r$ with equation (3).

Theorem 1. *[S98, Corollary 1.4.19]. Assume a pair of square matrices A (non-singular) and E such that $||A^{-1}E|| \le 1$. Then $||(A + E)^{-1}|| \le \frac{||A^{-1}||}{1-||A^{-1}E||}$ and moreover $\frac{||(A+E)^{-1}-A^{-1}||}{||A^{-1}||} \le \frac{||A^{-1}||}{1-||A^{-1}E||}$.*

2.3 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 the ratio $\sigma_1(A)/\sigma_\rho(A)$ is large. If the ratio is reasonably bounded, then the matrix is *well conditioned*. An $m \times n$ matrix A has a *numerical rank* $r = \mathrm{nrank}(A) \le \rho = \mathrm{rank}(A)$ if the ratios $\sigma_j(A)/||A||$ are small for $j > r$ but not for $j \le r$.

Remark 1. One can specify the adjective "small" above as "smaller than a fixed positive tolerance" and similarly specify "closely" and "well conditioned". The specification can be a challenge, e.g., for the matrix $\mathrm{diag}(1.1^{-j})_{j=0}^{999}$.

If a well conditioned $m \times n$ matrix A has a rank $\rho < l = \min\{m, n\}$, then all its close neighbors have numerical rank ρ and almost all of them have rank l. Conversely, if a matrix A has a positive numerical rank r, then the r-truncation A_r, obtained by setting to 0 all singular values $\sigma_j(A)$ for $j > r$, is a well conditioned rank-r approximation to the matrix A within the error norm bound $\sigma_{r+1}(A)$ (cf. (2)). It follows that a matrix is ill conditioned if and only if it is close to a matrix having a smaller rank and a matrix has a numerical rank r if and only if it can be closely approximated by a well conditioned matrix having full rank r. Rank-revealing factorizations of a matrix A that has a small numerical rank r, but possibly has a large rank ρ, produce its rank-r approximations at a lower computational cost [P00]. The randomized algorithms of [HMT11] decrease the computational cost further. An $m \times n$ matrix has *generic conditioning profile* if it has a numerical rank r and if its leading $i \times i$ blocks are nonsingular and well conditioned for $i = 1, \dots, r$. Such matrix is *strongly well conditioned* if it has full numerical rank $r = \min\{m, n\}$.

2.4 Block Gaussian elimination and GENP

For a nonsingular 2×2 block matrix $A = \begin{pmatrix} B & C \\ D & E \end{pmatrix}$ with a nonsingular *pivot block* $B = A^{(k)}$, define $S = S(A^{(k)}, A) = E - DB^{-1}C$, the *Schur complement* of $A^{(k)}$ in A, and the block factorization,

$$A = \begin{pmatrix} I_k & O_{k,r} \\ DB^{-1} & I_r \end{pmatrix} \begin{pmatrix} B & O_{k,r} \\ O_{r,k} & S \end{pmatrix} \begin{pmatrix} I_k & B^{-1}C \\ O_{k,r} & I_r \end{pmatrix}. \tag{4}$$

Apply this factorization recursively to the pivot block B and its Schur complement S and arrive at the block Gaussian elimination process, completely defined by the sizes of the pivot blocks. The recursive process either fails, where its pivot block turns out to be singular, in particular where it is a vanishing pivot entry of GENP, or can continue until all pivot blocks become nonzero scalars. When this occurs we arrive at GENP. Factorization (4) defines the block elimination of the first k columns of the matrix A, whereas $S = S(A^{(k)}, A)$ is the matrix produced at this elimination step. Now assume that the pivot dimensions $d_1, \dots, d_r$ and $\bar{d}_1, \dots, \bar{d}_{\bar{r}}$ of two block elimination processes sum to the same integer k, that is $k = d_1 + \cdots + d_r = \bar{d}_1 + \cdots + \bar{d}_{\bar{r}}$. Then verify that both processes produce the same Schur complement $S = S(A^{(k)}, A)$.

Theorem 2. *In every step of the recursive block factorization process based on (4) every diagonal block of a block diagonal factor is either a leading block of the input matrix A or the Schur complement $S(A^{(h)}, A^{(k)})$ for some integers h and k such that $0 < h < k \le n$ and $S(A^{(h)}, A^{(k)}) = (S(A^{(h)}, A))^{(h)}$.*

Corollary 2. *The recursive block factorization process based on equation (4) can be completed by involving no vanishing pivot elements and no singular pivot blocks if and only if the input matrix A has generic rank profile.*

Proof. Combine Theorem 2 with equation $\det A = (\det B) \det S$, implied by (4).

The following result shows that the pivot blocks are not close to singular matrices where block Gaussian elimination (e.g., GENP) is applied to a strongly well conditioned input matrix.

Theorem 3. *(Cf. [PQZ13, Theorem 5.1].) Assume GENP or block Gaussian elimination applied to an $n \times n$ matrix A and write $N = ||A||$ and $N_- = \max_{j=1}^{n} ||(A^{(j)})^{-1}||$, and so $N_- N \geq ||A||\ ||A^{-1}|| \geq 1$. Then the absolute values of all pivot elements of GENP and the norms of all pivot blocks of block Gaussian elimination do not exceed $N_+ = N + N_- N^2$, whereas the absolute values of the reciprocals of these elements and the norms of the inverses of the blocks do not exceed N_-.*

Proof. Observe that the inverse S^{-1} of the Schur complement S in (4) is the southeastern block of the inverse A^{-1} and obtain $||B|| \leq N$, $||B^{-1}|| \leq N_-$, and $||S^{-1}|| \leq ||A^{-1}|| \leq N_-$. Moreover $||S|| \leq N + N_- N^2$, due to (4). Now the claimed bound follows from Theorem 2.

Invert (4) to obtain $A^{-1} = \begin{pmatrix} I_k & -B^{-1}C \\ O_{k,r} & I_r \end{pmatrix} \begin{pmatrix} B^{-1} & O_{k,r} \\ O_{r,k} & S^{-1} \end{pmatrix} \begin{pmatrix} I_k & O_{k,r} \\ -DB^{-1} & I_r \end{pmatrix}$, extend this factorization recursively, apply it to the inversion of the matrix A and the solution of a linear system $A\mathbf{x} = \mathbf{b}$, and extend the above analysis.

Remark 2. For a strongly nonsingular input matrix A block factorization (4) can be extended to computing the complete recursive factorization, which defines GENP. By virtue of Theorem 3 the norms of the inverses of all pivot blocks involved in this computation are at most N_-. If the matrix A is also strongly well conditioned, then we have a reasonable upper bound on N_-, and so in view of Theorem 1 the inversion of all pivot blocks is numerically safe. In this case we say that *GENP and block Gaussian elimination are locally safe* for the matrix A. In a sense locally safe GENP and block Gaussian elimination are at least as safe numerically as GEPP because a locally safe recursive factorization involves neither divisions by absolutely small pivot entries (recall that pivoting has been introduced precisely in order to avoid such divisions) nor inversions of ill conditioned pivot blocks. Let us also compare the magnification of the perturbation norm bounds of Theorem 1 in GEPP and in the process of recursive factorization, which defines GENP and block Gaussian elimination. We observe immediately that in the recursive factorization only the factors of the leading blocks and the Schur complements can contribute to this magnification, namely at most $\log_2(n)$ such factors can contribute to the norm of each of the output triangular or block triangular factors L and U. This implies the upper bound $(N_+ N_-)^{\log_2(n)}$ on their norms, which can be compared favorably to the sharp upper bound on the growth factor 2^{n-1} for GEPP [GL96, page 119] and [S98, Theorem 3.4.12].

3 Singular values of the matrix products (deterministic estimates) and GENP with preprocessing

Fact 1 reduces the tasks of inverting a nonsingular and well conditioned matrix A and solving a linear system $A\mathbf{x} = \mathbf{b}$ to similar tasks for the matrix FAH and multipliers F and H of our choice. Remark 2 motivates the choice for which the matrix FAH is strongly nonsingular and strongly well conditioned. In Section 4.2 we prove that this is likely to occur already where one of the multipliers F and H is the identity matrix I and another one is Gaussian matrix, and therefore also where both F an H are independent Gaussian matrices. In this section we prepare background for that proof by estimating the norms of the matrices $(FA)_{k,k} = F_{k,m}A_{m,k}$ and $(AH)_{k,k} = A_{k,n}H_{n,k}$ and of their inverses for general (possibly nonrandom) multipliers F and H. We will keep writing M_s^T and M_s^+ for $(M_s)^T$ and $(M_s)^+$, respectively, where M can stand for A, F, or H, say, and where s can be any subscript such as a pair (k, l) or (A, r). We begin with two simple lemmas.

Lemma 1. *Suppose S and T are square orthogonal matrices. Then $\sigma_j(SA) = \sigma_j(AT) = \sigma_j(A)$ for all j.*

Lemma 2. *Suppose $\Sigma = \mathrm{diag}(\sigma_i)_{i=1}^n$, $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n$, $F \in \mathbb{R}^{r\times n}$, $H \in \mathbb{R}^{n\times r}$. Then $\sigma_j(F\Sigma) \geq \sigma_j(F)\sigma_n$ and $\sigma_j(\Sigma H) \geq \sigma_j(H)\sigma_n$ for all j. If also $\sigma_n > 0$, then $\mathrm{rank}(F\Sigma) = \mathrm{rank}(F)$ and $\mathrm{rank}(\Sigma H) = \mathrm{rank}(H)$.*

The following theorem bounds the norms $||(FA)^+||$ and $||(AH)^+||$ for three matrices A, F and H.

Theorem 4. *Suppose $A \in \mathbb{R}^{m\times n}$, $F \in \mathbb{R}^{r\times m}$, $H \in \mathbb{R}^{n\times r}$, $r \leq \rho$ for $\rho = \mathrm{rank}(A)$, $A = S_A\Sigma_A T_A^T$ (cf. (1)), $\widehat{F} = FS_A$, and $\widehat{H} = T_A^T H$. Then*

$$\sigma_j(FA) \geq \sigma_k(A)\ \sigma_j(\widehat{F}_{r,k}) \text{ for all } k \leq m \text{ and all } j, \tag{5}$$

$$\sigma_j(AH) \geq \sigma_l(A)\ \sigma_j(\widehat{H}_{l,r}) \text{ for all } l \leq n \text{ and all } j. \tag{6}$$

Proof. Note that $AH = S_A\Sigma_A T_A^T H$, and so $\sigma_j(AH) = \sigma_j(\Sigma_A T_A^T H) = \sigma_j(\Sigma_A\widehat{H})$ for all j by virtue of Lemma 1, because S_A is a square orthogonal matrix. Furthermore it follows from Fact 2 that $\sigma_j(\Sigma_A\widehat{H}) \geq \sigma_j(\Sigma_{l,A}\widehat{H}_{l,r})$ for all $l \leq n$. Combine this bound with the latter equations, then apply Lemma 2, and obtain bound (6). Similarly deduce bound (5).

Corollary 3. *Keep the assumptions and definitions of Theorem 4. Then*
(i) $\sigma_r(AH) \geq \sigma_\rho(A)\sigma_r(\widehat{H}_{\rho,r}) = \sigma_r(\widehat{H}_{\rho,r})/||A^+||$,
(ii) $||(AH)^+|| \leq ||A^+||\ ||\widehat{H}_{\rho,r}^+||$ *if* $\mathrm{rank}(AH) = \mathrm{rank}(\widehat{H}_{\rho,r}) = r$,
(iii) $\sigma_r(FA) \geq \sigma_\rho(A)\sigma_r(\widehat{F}_{r,\rho}) = \sigma_r(\widehat{F}_{r,\rho})/||A^+||$, *and*
(iv) $||(FA)^+|| \leq ||A^+||\ ||\widehat{F}_{r,\rho}^+||$ *if* $\mathrm{rank}(FA) = \mathrm{rank}(\widehat{F}_{r,\rho}) = r$.

Proof. Substitute $j = r$ and $l = \rho$ into bound (6), recall (3), and obtain part (i). If $\mathrm{rank}(AH) = \mathrm{rank}(\widehat{H}_{l,r}) = r$, then apply (3) to obtain that $\sigma_r(AH) = 1/||(AH)^+||$ and $\sigma_r(\widehat{H}_{l,r}) = 1/||\widehat{H}^+_{l,r}||$. Substitute these equations into part (i) and obtain part (ii). Similarly prove parts (iii) and (iv).

Let us extends Theorem 4 to the leading blocks of the matrix products.

Corollary 4. *Keep the assumptions and definitions of Theorem 4 and fix two positive integers k and l such that $k \le m$, $l \le n$. Then (i) $||(FA)^+_{k,l}|| \le ||\widehat{F}^+_{k,m}||\ ||A^+_{m,l}|| \le ||\widehat{F}^+_{k,m}||\ ||A^+||$ if $m \ge n = \rho$ and if the matrices $(FA)_{k,l}$ and $\widehat{F}_{k,m}$ have full rank, whereas (ii) $||(AH)^+_{k,l}|| \le ||\widehat{H}^+_{n,l}||\ ||A^+_{k,n}|| \le ||\widehat{H}^+_{n,l}||\ ||A^+||$ if $m = \rho \le n$ and if the matrices $(AH)_{k,l}$ and $\widehat{H}_{n,l}$ have full rank.*

Proof. Recall that $(FA)_{k,l} = F_{k,m}A_{m,l}$ and the matrix $A_{m,l}$ has full rank if $m \ge n = \rho$. Apply Corollary 3 for A and F replaced by $A_{m,l}$ and $F_{k,m}$, respectively, and obtain that $||(FA)^+_{k,l}|| \le ||\widehat{F}^+_{k,m}||\ ||A^+_{m,l}||$. Combine (3) and Corollary 1 and deduce that $||A^+_{m,l}|| \le ||A^+||$. Combine the two latter inequalities to complete the proof of part (i). Similarly prove part (ii).

The following definition formalizes the assumptions of Corollaries 3 and 4.

Definition 1. *Assume the matrices A, F, $\widehat{F}$, H, and $\widehat{H}$ of Theorem 4. Then the matrix pair (A, H) (resp. (F, A)) has* full rank *if the matrices AH, $\widehat{H}_{\rho,r}$ (resp. FA, $\widehat{F}_{r,\rho}$)) and A have full rank. This pair has full rank and is* well conditioned *if in addition the matrices $\widehat{H}_{\rho,r}$ (resp. $\widehat{F}_{r,\rho}$)) and A are well conditioned, whereas it has* generic rank profile *if $\mathrm{rank}(A) = \rho$ and $\mathrm{rank}((AH)_{k,k}) = \mathrm{rank}(\widehat{H}_{\rho,k}) = k$ (resp. $\mathrm{rank}((FA)_{k,k}) = \mathrm{rank}(\widehat{F}_{k,\rho}) = k$) for $k = 1, \dots, r$. The pair has generic rank profile and is* strongly well conditioned *if in addition the matrices $\widehat{H}_{\rho,k}$ (resp. $\widehat{F}_{k,\rho}$) for $k = 1, \dots, r$ are well conditioned.*

Remark 3. Fact 1, Corollary 2 and Theorem 3 together imply the following guiding rule. Suppose $A \in \mathbb{R}^{m\times n}$, $F \in \mathbb{R}^{r\times m}$, $H \in \mathbb{R}^{n\times r}$, $r \le \rho = \mathrm{rank}(A)$, and the matrix pair (A, H) for $m \le n$ or (F, A) for $m \ge n$ has generic rank profile and is strongly well conditioned. Then GENP is locally safe for the matrix products AH or FA, respectively (see Remark 2 on the concept "locally safe").

4 Benefits of using random matrix multipliers

Next we define Gaussian matrices and recall their basic properties. In Section 4.2 we show that these multipliers are expected to support locally safe GENP, and in Section 4.3 comment on using non-Gaussian random multipliers.

4.1 Gaussian matrix, its rank, norm and condition estimates

Definition 2. *A matrix is* standard Gaussian random *or just* Gaussian *if it is filled with i.i.d. Gaussian random variables having mean* 0 *and variance* 1.

Fact 4. *A Gaussian matrix is rank deficient with probability 0.*

Proof. Assume a rank deficient $m \times n$ matrix where $m \geq n$, say. Then the determinants of all its $n \times n$ submatrices vanish. This implies $\binom{m}{n}$ polynomial equations on the entries, that is the rank deficient matrices form an algebraic variety of a lower dimension in the linear space $\mathbb{R}^{m \times n}$. (V is an algebraic variety of a dimension $d \leq N$ in the space $\mathbb{R}^N$ if it is defined by $N - d$ polynomial equations and cannot be defined by fewer equations.) Clearly such a variety has Lebesgue (uniform) and Gaussian measure 0, both being absolutely continuous with respect to one another.

Corollary 5. *A Gaussian matrix has generic rank profile with probability 1.*

Hereafter $\nu_{j,m,n}$ denote the random variables $\sigma_j(G)$ for Gaussian $m \times n$ matrix G and all j, whereas $\nu_{m,n}$, $\nu^+_{m,n}$, and $\kappa_{m,n}$ denote the random variables $||G||$, $||G^+||$, and $\kappa(G) = ||G||\ ||G^+||$, respectively. Note that $\nu_{j,n,m} = \nu_{j,m,n}$, $\nu_{n,m} = \nu_{m,n}$, $\nu^+_{n,m} = \nu^+_{m,n}$, and $\kappa_{n,m} = \kappa_{m,n}$.

Theorem 5. *(Cf. [DS01, Theorem II.7].) Suppose* $h = \max\{m, n\}$, $t \geq 0$, *and* $z \geq 2\sqrt{h}$. *Then*
Probability$\{\nu_{m,n} > z\} \leq \exp(-(z - 2\sqrt{h})^2/2\}$ *and*
Probability$\{\nu_{m,n} > t + \sqrt{m} + \sqrt{n}\} \leq \exp(-t^2/2)$.

Theorem 6. *Suppose* $m \geq n$, *and* $x > 0$ *and write* $\Gamma(x) = \int_0^\infty \exp(-t)t^{x-1}dt$ *and* $\zeta(t) = t^{m-1}m^{m/2}2^{(2-m)/2}\exp(-mt^2/2)/\Gamma(m/2)$. *Then*
(i) Probability $\{\nu^+_{m,n} \geq m/x^2\} < \frac{x^{m-n+1}}{\Gamma(m-n+2)}$ *for* $n \geq 2$ *and*
(ii) Probability $\{\nu^+_{m,1} \geq x\} \leq (m/2)^{(m-2)/2}/(\Gamma(m/2)x^m)$.

Proof. (i) See [CD05, Proof of Lemma 4.1]. (ii) $G \in \mathbb{R}^{m \times 1}$ is a vector of length m. So, with probability 1, $G \neq 0$, $\text{rank}(G) = 1$, $||G^+|| = 1/||G||$, and consequently Probability $\{||G^+|| \geq x\}$ = Probability $\{||G|| \leq 1/x\} = \int_0^{1/x} \zeta(t)dt$. Note that $\exp(-mt^2/2) \leq 1$, and so $\int_0^{1/x} \zeta(t)dt < c_m \int_0^{1/x} t^{m-1}dt = c_m/(mx^m)$ where $c_m = m^{m/2}2^{(2-m)/2}/\Gamma(m/2)$.

The following condition estimates from [CD05, Theorem 4.5] are quite tight for large values x, but for $n \geq 2$ even tighter estimates (although more involved) can be found in [ES05]. (See [D88] and [E88] on the early study.)

Theorem 7. *With probability 1, we have* $\kappa_{m,1} = 1$. *If* $m \geq n \geq 2$, *then*
Probability$\{\kappa_{m,n}m/(m-n+1) > x\} \leq \frac{1}{2\pi}(6.414/x)^{m-n+1}$ *for* $x \geq m-n+1$.

Corollary 6. *A Gaussian matrix has generic rank profile with probability 1 and is expected to be well conditioned.*

Proof. Combine Corollary 5 and Theorem 7.

4.2 Supporting GENP with Gaussian multipliers

Lemma 3. *Suppose H is Gaussian matrix, S and T are orthogonal matrices, $H \in \mathbb{R}^{m\times n}$, $S \in \mathbb{R}^{k\times m}$, and $T \in \mathbb{R}^{n\times k}$ for some k, m, and n. Then SH and HT are Gaussian matrices.*

Theorem 8. *Suppose $A \in \mathbb{R}^{m\times n}$, $F \in \mathbb{R}^{r\times m}$, $H \in \mathbb{R}^{m\times r}$, F and H are Gaussian matrices, and* $\mathrm{rank}(A) = \rho$. *Then* $\mathrm{rank}(FA) = \mathrm{rank}(AH) = \min\{r, \rho\}$ *with probability 1.*

Proof. Suppose $A = S_A \Sigma_A T_A^T$ is SVD of (1). Then $FA = FS_A\Sigma_A T_A^T = G\Sigma_A T_A^T$ where $G = FS_A$ is Gaussian $r \times m$ matrix by virtue of Lemma 3. Clearly $\mathrm{rank}(FA) = \mathrm{rank}(G\Sigma_A T_A^T) = \mathrm{rank}(G\Sigma_A)$ because T_A is a square orthogonal matrix. Moreover $\mathrm{rank}(G\Sigma_A) = \mathrm{rank}(GD_\rho)$ where $D_\rho = \mathrm{diag}(I_\rho, O_{m-\rho,n-\rho})$, and so GD_ρ is Gaussian $r \times \rho$ matrix because it is a submatrix of the Gaussian matrix G. Therefore $\mathrm{rank}(FA) = \mathrm{rank}(GD_\rho)$ is equal to $\min\{r, \rho\}$ with probability 1 by virtue of Fact 4. Similarly obtain that $\mathrm{rank}(AH) = \min\{r, \rho\}$ with probability 1.

Corollary 7. *Keep the assumptions and definitions of Theorem 4. Suppose the matrix A has full rank $\rho = \min\{m, n\}$, $k \le r \le \rho$, and $F = FS_A$ and $H = T_A^T H$ are Gaussian matrices. Then (i) so are the matrices $\widehat{F}$, $\widehat{H}$ and all their submatrices, in particular $\widehat{F}_{k,\rho}$ and $\widehat{H}_{\rho,k}$, and (ii) with probability* 1, $\mathrm{rank}((AH)_{k,k}) = k$ *if $m \le n$,* $\mathrm{rank}((FA)_{k,k}) = k$ *if $m \le n$, and* $\mathrm{rank}(\widehat{H}_{\rho,k}) = \mathrm{rank}(\widehat{F}_{k,\rho}) = k$.

Proof. If H and F are Gaussian matrices, then so are the matrices $\widehat{H}$ and $\widehat{F}$ by virtue of Lemma 3. Consequently so are all their submatrices. This proves parts (i) and by virtue of Fact 4 also implies the equations $\mathrm{rank}(\widehat{H}_{\rho,k}) = \mathrm{rank}(\widehat{F}_{k,\rho}) = k$ of part (ii). Now recall that $(AH)_{k,k} = A_{k,n}H_{n,k}$, and hence $\mathrm{rank}((AH)_{k,k}) = \mathrm{rank}(A_{k,n}H_{n,k})$. This is equal to $\mathrm{rank}(A_{k,n})$ with probability 1 by virtue of Theorem 8 because $H_{n,k}$ is a Gaussian matrix and because $k \le \rho \le n$. Finally obtain that $\mathrm{rank}(A_{k,n}) = k$ for $k \le \rho = m$, and so $\mathrm{rank}((AH)_{k,k}) = k$. Similarly prove that $\mathrm{rank}((FA)_{k,k}) = k$ for $k \le \rho = n$.

Corollary 8. *The choice of Gaussian multipliers F where $m \le n$ or H where $m \ge n$ is expected to satisfy the assumptions of Remark 3 (thus supporting application of GENP to the matrix FA where $m \le n$ or AH where $m \ge n$) provided that the $m \times n$ matrix A is nonsingular and well conditioned.*

Proof. Combine Corollaries 6 and 7.

4.3 Structured random multipliers

Given $n \times n$ matrices A, F and H, we need $2n^3 - n^2$ flops to compute each of the products FA and AH, but we only need order of $n^2 \log(n)$ flops to compute such products where F and H are circulant matrices (cf. [P01]). Furthermore we need just n random parameters to define a Gaussian circulant $n \times n$ matrix

$C = (c_{i-j \mod n})_{i,j=0}^{n-1}$, expected to be very well conditioned [PQa]. Can we extend our results to these random multipliers? The proof and consequently the claim of Fact 4 can be immediately extended, and so we can still satisfy with probability 1 the generic rank profile assumption of Remark 3. Moreover we can satisfy that assumption with probability close to 1 even where we fill the multipliers F and H with i.i.d. random variables defined under the discrete uniform probability distribution over a fixed sufficiently large finite set (see Appendix B). We cannot extend our proof that the pair (T_A, H) is strongly well conditioned, however, because we cannot extend Lemma 3 to this case. Nevertheless empirically circulant random multipliers turn out to support GENP quite strongly (see our next section and also compare [HMT11], and [M11] on randomized low-rank approximation of a matrix). By engaging simultaneously two independent structured random multipliers F and H we may enhance their power.

5 Numerical Experiments

Our numerical experiments with general, Hankel, Toeplitz and circulant random 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 assuming the standard Gaussian probability distribution. The tests have been designed by the first author and performed by his coauthors.

Tables 3 and 4 show the results of our tests of the solution of nonsingular well conditioned linear systems $A\mathbf{y} = \mathbf{b}$ of n equations with random normalized vectors $\mathbf{b}$ whose coefficient matrices had singular $n/2 \times n/2$ leading principal blocks for $n = 64, 256, 1024$. Namely by following [H02, Section 28.3] we generated the $n \times n$ input matrices $A = \begin{pmatrix} A_k & B \\ C & D \end{pmatrix}$ where A_k was a $k \times k$ matrix, B, C and D were Toeplitz random matrices such that $||B|| \approx ||C|| \approx ||D|| \approx ||A_k|| \approx 1$, $n = 2^s$, $s = 6, 7, 8, 10$, and $A_k = U\Sigma V^T$ for $k = n/2$, $\Sigma = \mathrm{diag}(\sigma_i)_{i=1}^{k}$, $\sigma_i = 1$ for $i = 1, \ldots, k-h$, $\sigma_i = 0$ for $i = k-h+1, \ldots, k$, $h = 4$, and U and V were $k \times k$ orthonormal random matrices, computed as the unique $k \times k$ factors $Q(V)$ in the QR factorization of $k \times k$ random matrices V. We have performed 100 numerical tests for each dimension n and have computed the maximum, minimum and average relative residual norms $||A\mathbf{y} - \mathbf{b}||/||\mathbf{b}||$ as well as the standard deviation.

In our tests the norms $||A^{-1}||$ ranged from 70 to 4×10^6 (see Table 1), and so GEPP was expected to output accurate solutions to the linear systems $A\mathbf{y} = \mathbf{b}$, and indeed we always observed this in our tests (see Table 2). GENP, however, was expected to fail for these systems, because the leading block A_k of the matrix A was singular, having nullity $k - \mathrm{rank}(A_k) = 4$. Indeed this has caused poor performance of GENP in our tests, which has consistently output corrupted solutions, with the residual norms ranging from 10 to 10^8. In view

of Remark 8 we expected to fix this deficiency by means of multiplication by Gaussian random matrices, and indeed in our tests we observed consistently the residual norms below 4×10^{-9} for all inputs (see Table 3). Furthermore the tests showed the same power of preconditioning where we used the Gaussian circulant multipliers (see Table 4). As could be expected the output accuracy of GENP preprocessed with nonunitary random multipliers tended to deteriorate a little versus GEPP in our tests, but the output residual norms were small enough to support application of the inexpensive iterative refinement, whose single step decreased the output residual norm by factors of 10^h for $h > 1$ in the case of Gaussian multipliers and $h \geq 2$ in the case of Gaussian circulant multipliers (see Tables 3 and 4 and Remark 2).

Table 1. the norms $||A||^{-1}$ for the input matrices A

dimension	**min**	**max**	**mean**	**std**
64	6.9×10^{1}	2.3×10^{3}	4.6×10^{4}	6.4×10^{3}
256	8.4×10^{2}	1.1×10^{4}	5.8×10^{5}	5.8×10^{4}
1024	3.5×10^{3}	9.9×10^{4}	4.0×10^{6}	4.3×10^{5}

Table 2. Relative residual norms of GEPP

dimension	**min**	**max**	**mean**	**std**
64	2.0×10^{-15}	6.9×10^{-13}	3.2×10^{-14}	8.9×10^{-14}
256	1.4×10^{-14}	1.3×10^{-12}	1.2×10^{-13}	1.9×10^{-13}
512	5.2×10^{-14}	4.6×10^{-11}	1.0×10^{-12}	4.9×10^{-12}
1024	1.2×10^{-13}	1.0×10^{-09}	1.2×10^{-11}	1.0×10^{-10}

Table 3. Relative residual norms: GENP with Gaussian random multipliers

dimension	**iterations**	**min**	**max**	**mean**	**std**
64	0	3.6×10^{-13}	1.4×10^{-11}	5.3×10^{-12}	8.4×10^{-12}
64	1	7.3×10^{-15}	2.8×10^{-13}	3.2×10^{-14}	6.4×10^{-14}
256	0	5.1×10^{-12}	3.8×10^{-9}	9.2×10^{-10}	4.7×10^{-11}
256	1	4.8×10^{-15}	9.2×10^{-10}	8.6×10^{-11}	2.1×10^{-11}

Table 4. Relative residual norms: GENP with Gaussian circulant random multipliers

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}

6 Conclusions

It is well known that Gaussian (that is standard Gaussian random) matrices tend to be well conditioned, and this property motivates our preprocessing of well conditioned nonsingular input matrices with Gaussian multipliers to support the application of GENP and block Gaussian elimination. Both of these algorithms can fail in practical numerical computations without preprocessing, but we prove that with Gaussian multipliers the algorithms are expected to be locally safe, which would achieve the purpose of pivoting. Namely the absolute values of the reciprocals of all pivot elements of GENP and the norms of the inverses of all pivot blocks of block Gaussian elimination are likely to be reasonably bounded. Our tests were in good accordance with that formal study. We generated matrices that were hard for GENP, but the problems were consistently avoided where we preprocessed the inputs with Gaussian multipliers. Moreover empirically we obtained essentially the same results even where we used circulant random multipliers. Their choice accelerates multiplication significantly, and particularly dramatically in the important case where the input matrix has Toeplitz structure. That choice also limits randomization to n random parameters for an $n \times n$ input. Formal support for the empirical power of these multipliers is a natural research challenge.

Acknowledgements: Our research has been supported by NSF Grant CCF–1116736 and PSC CUNY Awards 64512–0042 and 65792–0043.

References

[CD05] Z. Chen, J. J. Dongarra, Condition Numbers of Gaussian Random Matrices, *SIAM. J. on Matrix Analysis and Applications*, **27**, 603–620, 2005.

[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.

[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.

[M11] M. W. Mahoney, Randomized Algorithms for Matrices and Data, *Foundations and Trends in Machine Learning*, NOW Publishers, **3, 2**, 2011. (Abridged version in: Advances in Machine Learning and Data Mining for Astronomy, edited by M. J. Way, et al., pp. 647-672, 2012.)

[P00] 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.

[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.

[PQa] V. Y. Pan, G. Qian, Estimating the Norms of Circulant and Toeplitz Random Matrices and Their Inverses, to appear.

[PQZ13] V. Y. Pan, G. Qian, A. Zheng, Randomized Preprocessing versus Pivoting, *Linear Algebra and Its Applications*, **438, 4**, 1883–1899, 2013.

[S80] J. T. Schwartz, Fast Probabilistic Algorithms for Verification of Polynomial Identities, *Journal of ACM*, **27**, **4**, 701–717, 1980.

[S95] J.-G. Sun, On Perturbation Bounds for QR Factorization, *Linear Algebra and Its Applications*, **215**, 95–111, 1995.

[S98] G. W. Stewart, *Matrix Algorithms, Vol I: Basic Decompositions*, SIAM, 1998.

[S03] G. W. Stewart, *Matrix Algorithms, Vol II: Eigensystems*, SIAM, 2003.

[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.

[Z79] R. E. Zippel, Probabilistic Algorithms for Sparse Polynomials, *Proceedings of EUROSAM'79, Lecture Notes in Computer Science*, **72**, 216–226, Springer, Berlin, 1979.

Appendix

A On the algebraic variety of low-rank matrices

The following simple result (not used in this paper) shows that the $m \times n$ matrices of a rank ρ form an algebraic variety of the dimension $d_\rho = (m + n - \rho)\rho$ in the space $\mathbb{R}^{m \times n}$, and clearly $d_\rho < mn$ for $\rho < \min\{m, n\}$.

Fact 5. *The set $\mathbb{A}$ of $m \times n$ matrices of rank ρ is an algebraic variety of dimension $(m+n-\rho)\rho$.*

Proof. Let A be an $m \times n$ matrix of a rank ρ with a nonsingular leading $\rho \times \rho$ block B and write $A = \begin{pmatrix} B & C \\ D & E \end{pmatrix}$. Then the $(m-\rho) \times (n-\rho)$ *Schur complement* $E - DB^{-1}C$ must vanish, which imposes $(m-\rho)(n-\rho)$ algebraic equations on the entries of the matrix A. Similar argument can be applied where any $\rho \times \rho$ submatrix of the matrix A (among $\binom{m}{\rho}\binom{n}{\rho}$ such submatrices) is nonsingular. Therefore $\dim \mathbb{A} = mn - (m-\rho)(n-\rho) = (m+n-\rho)\rho$.

B 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.

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 4. *[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 the set Δ^m. Then the polynomial vanishes in at most $d|\Delta|^{m-1}$ points of this set.*

Theorem 9. *Under the assumptions of Lemma 4 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 9. *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 $k \le l$ 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 9 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.