MANCHESTER 1824

The University of Manchester

M\cr NA
Manchester Numerical Analysis

Random Matrices Generating Large Growth in LU Factorization With Pivoting

Srikara Pranesh
Department of Mathematics
The University of Manchester

srikara.pranesh@manchester.ac.uk

23-07-2021

Joint work with

Nick Higham and Desmond Higham

Slides Available at – https://bit.ly/AN21_Sri

Introduction

Type	bits Signif(t)	Exp	Range	$u = 2^{-t}$
bfloat16	8	8	$10^{\pm 38}$	3.9×10^{-3}
fp16	11	5	$10^{\pm 5}$	4.9×10^{-4}
tf32	11	8	$10^{\pm 38}$	4.9×10^{-4}
fp32	24	8	$10^{\pm 38}$	6.0×10^{-8}
fp64	53	11	$10^{\pm 308}$	1.1×10^{-16}
fp128	113	45	$10^{\pm 4932}$	9.6×10^{-35}

Introduction

Type	bits Signif(t)	Exp	Range	$u = 2^{-t}$
bfloat16	8	8	$10^{\pm 38}$	3.9×10^{-3}
fp16	11	5	$10^{\pm 5}$	4.9×10^{-4}
tf32	11	8	$10^{\pm 38}$	4.9×10^{-4}
fp32	24	8	$10^{\pm 38}$	6.0×10^{-8}
fp64	53	11	$10^{\pm 308}$	1.1×10^{-16}
fp128	113	45	$10^{\pm 4932}$	9.6×10^{-35}

- Speed, energy, communication benefits.

Introduction

Type	bits Signif(t)	Exp	Range	$u = 2^{-t}$
bfloat16	8	8	$10^{\pm 38}$	3.9×10^{-3}
fp16	11	5	$10^{\pm 5}$	4.9×10^{-4}
tf32	11	8	$10^{\pm 38}$	4.9×10^{-4}
fp32	24	8	$10^{\pm 38}$	6.0×10^{-8}
fp64	53	11	$10^{\pm 308}$	1.1×10^{-16}
fp128	113	45	$10^{\pm 4932}$	9.6×10^{-35}

- Speed, energy, communication benefits.

▲ We need algorithms ....
- that use low precision.
- that are provably robust.

▲ **Mixed Precision Algorithms [Abdelfattah et.al, 2021].**

GMRES-IR [Carson, Higham 2018]

Given A and b in precision u, and $u_f \geq u \geq u_r$

solve $Ax_0 = b$ using the LU factors of precision $u_f > u$

- $r = b - Ax_0$, in $u_r < u$.
- Solve $\widetilde{A}d \equiv \widehat{U}^{-1}\widehat{L}^{-1}A = \widehat{U}^{-1}\widehat{L}^{-1}r$, at precision u using GMRES.
- Update $x_1 = \mathrm{fl}(x_0 + d)$ in precision u.

GMRES-IR [Carson, Higham 2018]

Given A and b in precision u, and $u_f \geq u \geq u_r$

solve $Ax_0 = b$ using the LU factors of precision $u_f > u$

- $r = b - Ax_0$, in $u_r < u$.
- Solve $\widetilde{A}d \equiv \widehat{U}^{-1}\widehat{L}^{-1}A = \widehat{U}^{-1}\widehat{L}^{-1}r$, at precision u using GMRES.
- Update $x_1 = \mathrm{fl}(x_0 + d)$ in precision u.

- For a GPU implementation Four times speedup, and 80% reduction in energy consumption. **[Haidar et.al, 2018]**
- Vendor and open source implementations available.
- Extension to SPD systems **[Higham, P, 2021]**, and least squares problems **[Carson, Higham, P, 2020]**.
- Algorithmic basis for the new **HPL-AI** benchmark.

Overflow in fp16 and its remedy

fp16 has narrow range:

	u	$x_{\min}^{s}$	$x_{\min}$	$x_{\max}$
fp16	4.88×10^{-4}	5.96×10^{-8}	6.10×10^{-5}	6.55×10^{4}

Can be remedied using two sided diagonal scaling for $Ax = b$. (Higham, **P**, Zounon 2019).

2DS Rounds $A \in \mathbb{R}^{n\times n}$ to the fp16 matrix $A^{(h)}$, scaling all elements to avoid overflow. $\theta \in (0, 1]$ is a parameter.

1: Apply any two-sided diagonal scaling algorithm to A, to obtain diagonal matrices R, S.
   `%squeeze entries between` $-1 \le (RAS)_{ij} \le 1$
2: $\mu = \theta x_{\max}$ `% make full use of the range`
3: $A^{(h)} = \mathrm{fl}_h(\mu(RAS))$

How to choose θ?

Choice of θ

- θ – headroom for further computation.
- In $PA = LU$,

$$|\ell_{ij}| \leq 1, \quad |u_{ij}| \leq \rho_n(A) \max_{ij} |a_{ij}|,$$

  $\rho_n(A)$ is the growth factor.
- In practice, ρ_n is **almost always small** (< 10) for partial pivoting. *Open problem to explain why!*
- **Usually** $\theta = 0.1$ is sufficient.

Choice of θ

- θ – headroom for further computation.
- In $PA = LU$,

$$|\ell_{ij}| \leq 1, \quad |u_{ij}| \leq \boldsymbol{\rho_n(A)} \max_{ij} |a_{ij}|,$$

  $\rho_n(A)$ is the growth factor.
- In practice, ρ_n is **almost always small** ($<$ **10**) for partial pivoting. *Open problem to explain why!*
- **Usually** $\theta = 0.1$ is sufficient.

- Rest of the talk –

When is $\theta = 0.1$ not sufficient?

The Growth Factor

Gaussian elimination on $A \in \mathbb{R}^{n\times n}$ produces $A = LU$.

With $A^{(1)} = A$, $A^{(n)} = U$, $A^{(k)} = \left(a_{ij}^{(k)}\right)$ matrix at kth stage of Gaussian elimination,

$$\rho_n(A) = \frac{\max_{i,j,k} |a_{ij}^{(k)}|}{\max_{i,j} |a_{ij}|} \geq 1.$$

The Growth Factor

Gaussian elimination on $A \in \mathbb{R}^{n \times n}$ produces $A = LU$.

With $A^{(1)} = A$, $A^{(n)} = U$, $A^{(k)} = \left(a_{ij}^{(k)}\right)$ matrix at kth stage of Gaussian elimination,

$$\rho_n(A) = \frac{\max_{i,j,k} |a_{ij}^{(k)}|}{\max_{i,j} |a_{ij}|} \geq 1.$$

Theorem (Wilkinson, 1961)

GE produces a computed solution $\widehat{x}$ to $Ax = b$ satisfying Then

$$(A + \Delta A)\widehat{x} = b, \qquad \|\Delta A\|_\infty \leq p(n)\rho_n u \|A\|_\infty,$$

where u is unit roundoff and p a low degree polynomial.

What We Know About the Growth Factor

- Without pivoting, ρ_n can arbitrarily large.
- With **partial pivoting**, $\rho_n \leq 2^{n-1}$, attained for

$$A_4 = \begin{bmatrix} 1 & 0 & 0 & 1 \\ -1 & 1 & 0 & 1 \\ -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & 1 \end{bmatrix}.$$

- **Wright (1993)** and **Foster (1994)** found applications where partial pivoting suffers exponential growth.
- **Higham & Higham (1989)** found orthogonal matrices with $\rho_n \gtrsim n/2$ for *any* pivoting strategy.

MATLAB Function

```
function g = gf(A)
%GF     Approximate growth factor.
%   g = GF(A) is an approximation to the
%   growth factor for LU factorization
%   with partial pivoting.
[~,U] = lu(A);
g = max(abs(U),[],'all')/max(abs(A),[],'all');
```

- This is a lower bound on $\rho_n(A)$.
- Can get exact growth factor using `gep.m` from **Matrix Computation Toolbox**.

```
>> rng(1); gf(randn(10))
ans =
   1.5088e+00
>> gf(randn(100))
ans =
   4.4874e+00
>> gf(randn(1000))
ans =
   1.5997e+01
>> gf(randn(10000))
ans =
   5.0505e+01
```

```
>> rng(1); gf(randn(10))
ans =
   1.5088e+00
>> gf(randn(100))
ans =
   4.4874e+00
>> gf(randn(1000))
ans =
   1.5997e+01
>> gf(randn(10000))
ans =
   5.0505e+01
>> gf(gallery('randsvd',1000,1e8,2,[],[],1))
ans =
   7.5329e+01
```

Randsvd Matrices (Mode 2)

$$A = P\Sigma Q^T \in \mathbb{R}^{n\times n}, \quad P^T P = Q^T Q = I,$$
$$\Sigma = \mathrm{diag}(1, \ldots, 1, \sigma_n), \quad 1 \geq \sigma_n \geq 0.$$

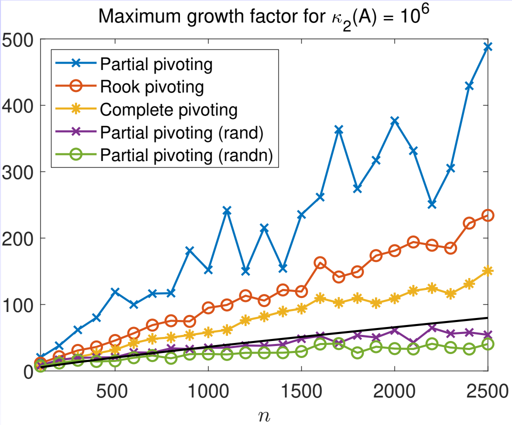

Randsvd Matrices

- P and Q are random orthogonal matrices drawn from Haar distribution.
- ```
  [Q,R] = qr(randn(n));
  Q = Q*diag(sign(diag(R)));
  ```
- More efficient method by **[Steward (1980)]**.
- **Demmel & McKenney (1989)**: LAPACK's test matrix generation suite.
- **H (1991, 1995)**: MATLAB function `randsvd`. Subsequently incorporated into `gallery('randsvd',...)`.
- **H & Zhang (2016)**: Matrix Depot package for Julia.
- Cheaper implementation proposed by **[Higham, & Fasi (2021)]**.
  ```

Growth Factor Lower Bound

Theorem (H & H, 1989)

Let $A \in \mathbb{C}^{n\times n}$ be nonsingular,

$$\alpha = \max_{i,j} |a_{ij}|, \quad \beta = \max_{i,j} \left| \left(A^{-1}\right)_{ij} \right|, \quad \theta = (\alpha\beta)^{-1}.$$

Then $\theta \le n$, and for any permutation matrices Π_r and Π_c such that $\Pi_r A \Pi_c$ has an LU factorization, the growth factor for GE without pivoting on $\Pi_r A \Pi_c$ satisfies

$$\rho_n(A) \ge \theta.$$

Growth for Random Orthogonal Matrices

Randsvd with $\sigma_n = 1$ gives $A = PQ^T$: **random orthogonal matrix from Haar distribution**.

Jiang (2005) shows that

$$\Pr\left(\max_{i,j} |a_{ij}| > 2\sqrt{\frac{\log(n)}{n}}(1+\epsilon)\right) \to 0$$

as $n \to \infty$ for any $\epsilon > 0$. Since $A^{-1} = A^T$, can take $\alpha = \beta = 2\sqrt{\log(n)/n}$ in the theorem and conclude

$$\rho_n(A) \gtrsim \frac{n}{4\log n}$$

for large n with high probability for *any* pivoting strategy.

Growth Factors for Random Orthogonal

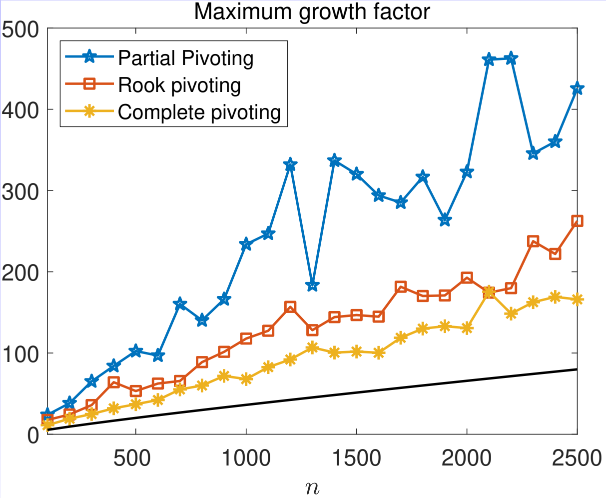

Proof of Large Growth for Randsvd

- The randsvd matrix is

$$A = PQ^T + (\sigma_n - 1)p_n q_n^T,$$

where p_n and q_n are the last columns of P and Q.

- If W is orthogonal and has large growth then a rank-1 perturbation of norm at most 1 *tends to preserve the large growth*.
- **Not particular** to W being Haar distributed.
- One approach is via **Sherman–Morrison formula**.

Direct Approach

Let W be orthogonal and

$$A = W + xy^T.$$

The U factor of W is given explicitly by

$$u_{ij} = \frac{\det\big(W(1\colon i, [1\colon i-1,\ j])\big)}{\det(W_{i-1})}, \quad i \le j,$$

where $W_j = W(1\colon j, 1\colon j)$. Find $\widetilde{U}$ factor of A satisfies

$$\frac{\widetilde{u}_{ij}}{u_{ij}} = \frac{1 + y([1\colon i-1,\ j])^T\ W(1\colon i, [1\colon i-1,\ j])^{-1}\ x(1\colon i)}{1 + y(1\colon i-1)^T\ W_{i-1}^{-1}\ x(1\colon i-1)}.$$

Effect of $O(n)$ Growth.

Effect of $O(n)$ Growth.

- $n = 10^7$ for today's largest dense $Ax = b$
  $\Rightarrow$ problems in single precision.
- For IEEE half precision and $\max_{i,j} |a_{ij}| = 1$, linear growth can cause overflow for $n = 7 \times 10^4$.
- For the previous example $\theta = 0.1$ wont be enough.
  (That's how these matrices were spotted.)

Numerical Experiments

Numerical Experiments

- mode-2 `randsvd` matrix with $\kappa_2(A) = 10^7$ and varying size.
- GMRES-IR and LU-IR are compared.
- (half,double,double) and (half, single, double) combinations are considered.
- `chop` by **[Higham, & P, 2019]** used for fp16 computation.

Results

		$\kappa_\infty(\widehat{U}^{-1}\widehat{L}^{-1}A)$	
n	ρ_n	(H, D, D)	(S, D, D)
500	53.29	3.30e+10	2.32e+03
750	103.80	3.62e+10	3.33e+03
1000	90.27	9.07e+10	4.84e+03
1250	102.03	1.57e+11	1.72e+04
1500	178.48	1.23e+11	1.51e+04
1750	186.22	2.11e+11	3.31e+04
2000	321.61	2.50e+11	1.48e+04
2250	349.27	3.84e+11	3.28e+04
2500	188.25	3.97e+11	1.34e+05

Results

	(H, D, D)		(S, D, D)	
n	LU-IR	GMRES-IR	LU-IR	GMRES-IR
500	–	3 (10)	12	3 (5)
750	–	3 (10)	–	3 (5)
1000	17	3 (13)	–	2 (4)
1250	–	4 (16)	16	3 (5)
1500	19	3 (12)	12	3 (5)
1750	–	3 (12)	19	3 (5)
2000	18	3 (12)	19	3 (5)
2250	–	2 (8)	–	3 (5)
2500	–	3 (13)	–	2 (4)

Summary

New class of **random**, dense $A \in \mathbb{R}^{n \times n}$ (*randsvd mode* 2) for which

- $\rho_n \gtrsim n/(4 \log n)$ for large n with **any form of pivoting**,
- $\kappa_2(A)$ can be **arbitrarily chosen**.

- GMRES-IR performs well despite lower quality LU factors.

Have been part of MATLAB **gallery** for many years but their growth properties had not been recognized.

D. J. Higham, N. J. Higham, and S. Pranesh. **Random matrices generating large growth in LU factorization with pivoting**. SIMAX, 2021.