





Sparse Matrices: Goals

- Perform standard matrix computations economically, i.e., without storing the zeros of the matrix.
- For typical Finite Element and Finite Difference matrices, the number of non-zero entries is $O(n)$



Sparse Matrices: MP example



Sparse Matrices

Some type of matrices contain many zeros.
Storing all those zero entries is wasteful!

How can we efficiently store large
matrices without storing tons of zeros?



- **Sparse matrices** (vague definition): matrix with few non-zero entries.
- For practical purposes: an $m \times n$ matrix is sparse if it has $O(\min(m, n))$ non-zero entries.
- This means roughly a constant number of non-zero entries per row and column.
- Another definition: “matrices that allow special techniques to take advantage of the large number of zero elements” (J. Wilkinson)

Sparse Matrices

EXAMPLE:

Number of operations required to add two square dense matrices:

$$O(n^2)$$

Number of operations required to add two sparse matrices **A** and **B**:

$$O(\text{nnz}(\mathbf{A}) + \text{nnz}(\mathbf{B}))$$

where $\text{nnz}(\mathbf{X})$ = number of non-zero elements of a matrix **X**

Popular Storage Structures

DNS	Dense	ELL	Ellpack-Itpack
BND	Linpack Banded	DIA	Diagonal
COO	Coordinate	BSR	Block Sparse Row
CSR	Compressed Sparse Row	SSK	Symmetric Skyline
CSC	Compressed Sparse Column	BSR	Nonsymmetric Skyline
MSR	Modified CSR	JAD	Jagged Diagonal
LIL	Linked List		

note: CSR = CRS, CCS = CSC, SSK = SKS in some references

We will focus on COO and CSR!

Dense (DNS)

$$A = \begin{bmatrix} 0. & 1.9 & 0. & -5.2 \\ 0.3 & 0. & 9.1 & 0. \\ 4.4 & 5.8 & 3.6 & 0. \\ 0. & 0. & 7.2 & 2.7 \end{bmatrix}$$

$A_{shape} = (nrow, ncol)$

$$A_{dense} = [0. \quad 1.9 \quad 0. \quad -5.2 \quad 0.3 \quad 0. \quad 9.1 \quad 0. \quad 4.4 \quad 5.8 \quad 3.6 \quad 0. \quad 0. \quad 0. \quad 7.2 \quad 2.7]$$

Row 0 Row 1 Row 2 Row 3

- Simple
- Row-wise
- Easy blocked formats
- Stores all the zeros

Coordinate (COO)

$$A = \begin{matrix} & \begin{matrix} 0 & 1 & 2 & 3 \end{matrix} \\ \begin{matrix} 0 \\ 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 0. & \underline{1.9} & 0. & \underline{-5.2} \\ \underline{0.3} & 0. & \underline{9.1} & 0. \\ \underline{4.4} & \underline{5.8} & \underline{3.6} & 0. \\ 0. & 0. & \underline{7.2} & \underline{2.7} \end{bmatrix} \end{matrix}$$

$$data = [-5.2 \ 0.3 \ 9.1 \ 1.9 \ 5.8 \ 7.2 \ 3.6 \ 4.4 \ 2.7]$$

$$data = [1.9 \ -5.2 \ 0.3 \ 9.1 \ 4.4 \ 5.8 \ 3.6 \ 7.2 \ 2.7]$$

$$row = [0 \ 0 \ 1 \ 1 \ 2 \ 2 \ 2 \ 3 \ 3]$$

$$col = [1 \ 3 \ 0 \ 2 \ 0 \ 1 \ 2 \ 2 \ 3]$$

- Simple
- Does not store the zero elements
- Not sorted
- *row* and *col*: array of integers
- *data*: array of doubles

$$row = [0 \ 1 \ 1 \ 0 \ 2 \ 3 \ 2 \ 2 \ 3]$$

$$col = [3 \ 0 \ 2 \ 1 \ 1 \ 2 \ 2 \ 0 \ 3]$$

Iclicker question

$$A = \begin{bmatrix} 1 & 0 & 0 & 2 & 0 \\ 3 & 4 & 0 & 5 & 0 \\ 6 & 0 & 7 & 8 & 9 \\ 0 & 0 & 10 & 11 & 0 \\ 0 & 0 & 0 & 0 & 12 \end{bmatrix}$$

```
data = [ 12.0  9.0  7.0  5.0  1.0  2.0  11.0  3.0  6.0  4.0  8.0  10.0 ]
row   = [  4    2    2    1    0    0    3    1    2    1    2    3    ]
col   = [  4    4    2    3    0    3    3    0    0    1    3    2    ]
```

How many integers are stored in COO format
(A has dimensions $n \times n$)?

A) nnz

B) n

C) $2 nnz$

D) n^2

E) $2 n$

row $\rightarrow nnz \rightarrow int$
col $\rightarrow nnz \rightarrow int$
data $\rightarrow nnz \rightarrow float$

Iclicker question

Representing a Sparse Matrix in Coordinate (COO) Form

1 point

Consider the following matrix:

$$A = \begin{bmatrix} 0 & 0 & 1.3 \\ -1.5 & 0.2 & 0 \\ 5 & 0 & 0 \\ 0 & 0.3 & 3 \\ 0 & 0 & 0 \end{bmatrix}$$

- A) 56 bytes
- B) 72 bytes
- C) 96 bytes
- D) 120 bytes
- E) 144 bytes

Suppose we store one row index (a 32-bit integer), one column index (a 32-bit integer), and one data value (a 64-bit float) for each non-zero entry in A . How many bytes in total are stored? Please note that 1 byte is equal to 8 bits.

$$\underbrace{2 \text{ nnz} * 32}_{\text{int}} + \underbrace{\text{nnz} * 64}_{\text{float}} = 12 * 32 + 6 * 64 = \frac{768}{8} = 96 \text{ bytes}$$

Compressed Sparse Row (CSR)

$$A = \begin{bmatrix} 1 & 0 & 0 & 2 & 0 \\ 3 & 4 & 0 & 5 & 0 \\ 6 & 0 & 7 & 8 & 9 \\ 0 & 0 & 10 & 11 & 0 \\ 0 & 0 & 0 & 0 & 12 \end{bmatrix}$$

0 → 2 non-zero
 1 → 3 "
 2 → 4
 3 → 2
 4 → 1

float (nnz)

row 0 row 1 row 2 row 3 row 4

data = [1.0 2.0 3.0 4.0 5.0 6.0 7.0 8.0 9.0 10.0 11.0 12.0]

int (nnz)

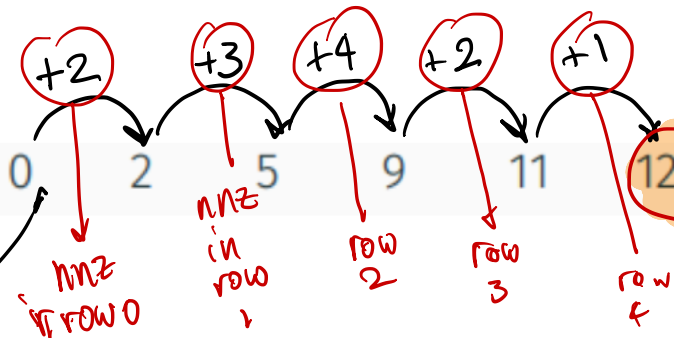
col = [0 3 0 1 3 0 2 3 4 2 3 4]

float (number of rows)

rowptr = [0 2 5 9 11 12]

total number of non-zero entries

always zero



Compressed Sparse Row (CSR)

$$A = \begin{bmatrix} 1 & 0 & 0 & 2 & 0 \\ 3 & 4 & 0 & 5 & 0 \\ 6 & 0 & 7 & 8 & 9 \\ 0 & 0 & 10 & 11 & 0 \\ 0 & 0 & 0 & 0 & 12 \end{bmatrix}$$

```
data    = [ 1.0  2.0  3.0  4.0  5.0  6.0  7.0  8.0  9.0 10.0 11.0 12.0 ]
col     = [ 0    3    0    1    3    0    2    3    4    2    3    4    ]
rowptr  = [ 0    2    5    9   11  12    ]
```

- Does not store the zero elements
- Fast arithmetic operations between sparse matrices, and fast matrix-vector product
- *col*: contain the column indices (array of *nnz* integers)
- *data*: contain the non-zero elements (array of *nnz* doubles)
- *rowptr*: contain the row offset (array of $n + 1$ integers)

Example - CSR format

$$A = \begin{bmatrix} 0 & 1 & 2 & 3 \\ 0. & 1.9 & 0. & -5.2 \\ 0. & 0. & 0. & 0. \\ 4.4 & 5.8 & 3.6 & 0. \\ 0. & 0. & 7.2 & 2.7 \end{bmatrix} \begin{matrix} \longrightarrow 2 \\ \longrightarrow 0 \\ \longrightarrow 3 \\ \longrightarrow 2 \end{matrix}$$

$$\text{data} = [1.9 \quad -5.2 \quad 4.4 \quad 5.8 \quad 3.6 \quad 7.2 \quad 2.7]$$

$$\text{col} = [1 \quad 3 \quad 0 \quad 1 \quad 2 \quad 2 \quad 3]$$


$$\text{rowptr} = [0 \quad 2 \quad 2 \quad 5 \quad 7]$$

Back to Image-Blur


★ How perturbation
"slowly" creeps in and
affects the perturbation
in the output

the shower faucet

how they are:



how they should be:



WHAT IT LOOKS LIKE



WHAT IT FEELS LIKE



Numerical experiments

Input has uncertainties:

- Errors due to representation with finite precision
- Error in the sampling

Once you select your numerical method , how much error should you expect to see in your **output**?

Is your method sensitive to errors (perturbation) in the input?

Demo “HilbertMatrix-ConditionNumber”

Hilbert Matrix

$$\begin{pmatrix} \underline{A} \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} = \begin{pmatrix} \underline{b} \end{pmatrix} \quad \underline{A} \underline{x} = \underline{b}$$

↳ this is the exact solution \underline{x}

→ generate random \underline{A} , use mat-vec to get \underline{b}

→ "undo" button

→ $\underline{A} \hat{\underline{x}} = \underline{b} \longrightarrow \hat{\underline{x}}$ is the approx solution (solve!)

→ error = $\|\hat{\underline{x}} - \underline{x}\|_2$

Sensitivity of Solutions of Linear Systems

Suppose we start with a non-singular system of linear equations $A \mathbf{x} = \mathbf{b}$.

We change the right-hand side vector \mathbf{b} (input) by a small amount $\Delta \mathbf{b}$.

How much the solution \mathbf{x} (output) changes, i.e., how large is $\Delta \mathbf{x}$?

$$\frac{\text{Output Relative error}}{\text{Input Relative error}} = \frac{\|\Delta \mathbf{x}\| / \|\mathbf{x}\|}{\|\Delta \mathbf{b}\| / \|\mathbf{b}\|} = \frac{\|\Delta \mathbf{x}\| \|\mathbf{b}\|}{\|\Delta \mathbf{b}\| \|\mathbf{x}\|}$$

Ax + AΔx = b + Δb

$$A \hat{\mathbf{x}} = \hat{\mathbf{b}} \rightarrow A \hat{\mathbf{x}} = A(\mathbf{x} + \Delta \mathbf{x}) = (\mathbf{b} + \Delta \mathbf{b}) \rightarrow A \Delta \mathbf{x} = \Delta \mathbf{b}$$

$$\frac{\text{Output Relative error}}{\text{Input Relative error}} = \frac{\overbrace{\|A^{-1} \Delta \mathbf{b}\|}^{\Delta \mathbf{x}} \underbrace{\|A \mathbf{x}\|}_{\mathbf{b}}}{\|\Delta \mathbf{b}\| \|\mathbf{x}\|} \leq \frac{\|A^{-1}\| \|\Delta \mathbf{b}\| \|A\| \|\mathbf{x}\|}{\|\Delta \mathbf{b}\| \|\mathbf{x}\|}$$

$$\frac{\|\Delta \mathbf{x}\|}{\|\mathbf{x}\|} \leq \|A^{-1}\| \|A\| \frac{\|\Delta \mathbf{b}\|}{\|\mathbf{b}\|}$$

↳ = \|A^{-1}\| \|A\|

Sensitivity of Solutions of Linear Systems

We can also add a perturbation to the matrix \mathbf{A} (input) by a small amount \mathbf{E} , such that

$$(\mathbf{A} + \mathbf{E}) \hat{\mathbf{x}} = \mathbf{b}$$

and in a similar way obtain:

$$\frac{\|\Delta \mathbf{x}\|}{\|\mathbf{x}\|} \leq \|\mathbf{A}^{-1}\| \|\mathbf{A}\| \frac{\|\mathbf{E}\|}{\|\mathbf{A}\|}$$

Condition number

Demo "HilbertMatrix-ConditionNumber"

get the condition number.

The condition number is a measure of sensitivity of solving a linear system of equations to variations in the input.

→ High cond(A) !
→ seems to correlate well

The condition number of a matrix A :

$$\text{cond}(A) = \|A^{-1}\| \|A\|$$

Recall that the induced matrix norm is given by

$$\|A\| = \max_{\|x\|=1} \|Ax\|$$

And since the condition number is relative to a given norm, we should be precise and for example write:

$$\text{cond}_2(A) \text{ or } \text{cond}_\infty(A)$$

Clicker question

$$\frac{\|\Delta \mathbf{x}\|}{\|\mathbf{x}\|} \leq \text{cond}(\mathbf{A}) \frac{\|\Delta \mathbf{b}\|}{\|\mathbf{b}\|}$$

Give an example of a matrix that is very well-conditioned (i.e., has a condition number that is good for computation). Select the best possible condition number(s) of a matrix?

~~A) $\text{cond}(\mathbf{A}) < 0$~~

~~B) $\text{cond}(\mathbf{A}) = 0$~~

~~C) $0 < \text{cond}(\mathbf{A}) < 1$~~

D) $\text{cond}(\mathbf{A}) = 1$

~~E) $\text{cond}(\mathbf{A}) = \text{large numbers}$~~

$\text{cond}(\mathbf{A}) = \|\mathbf{A}\| \|\mathbf{A}^{-1}\| > 0$

} why not these answers?

this will amplify perturbations in the output

Condition number

$$\frac{\|\Delta \mathbf{x}\|}{\|\mathbf{x}\|} \leq \underbrace{\text{cond}(\mathbf{A})}_{\text{small}} \frac{\|\Delta \mathbf{b}\|}{\|\mathbf{b}\|}$$

Small condition numbers mean not a lot of error amplification. Small condition numbers are good!

The identity matrix should be well-conditioned:

$$\|\mathbf{I}\| = \max_{\|\mathbf{x}\|=1} \|\mathbf{I} \mathbf{x}\| = 1$$

$$\text{cond}(\mathbf{I}) = \|\mathbf{I}\| \|\mathbf{I}^{-1}\| = 1$$

It turns out that this is the smallest possible condition number:

$$\text{cond}(\mathbf{A}) = \|\mathbf{A}^{-1}\| \|\mathbf{A}\| \geq \|\mathbf{A}^{-1} \mathbf{A}\| = \|\mathbf{I}\| = 1$$

$$\boxed{\text{cond}(\mathbf{A}) \geq 1}$$

If \mathbf{A}^{-1} does not exist, then $\text{cond}(\mathbf{A}) = \infty$ (by convention)

(SINGULAR)

Recall Induced Matrix Norms

$$\|A\|_1 = \max_j \sum_{i=1}^n |A_{ij}|$$


Maximum absolute column sum of the matrix A

$$\|A\|_\infty = \max_i \sum_{j=1}^n |A_{ij}|$$

Maximum absolute row sum of the matrix A

$$\|A\|_2 = \max_k \sigma_k$$

σ_k are the singular value of the matrix A



Clicker question

Condition Number of a Diagonal Matrix

What is the 2-norm-based condition number of the diagonal matrix

$$A = \begin{bmatrix} 100 & 0 & 0 \\ 0 & 13 & 0 \\ 0 & 0 & 0.5 \end{bmatrix}?$$

- A) 1
- B) 50
- C) 100
- D) 200

$$\|A\|_2 = \max_k \sigma_k = \max(100, 13, 0.5) = 100$$

$$\|A^{-1}\|_2 = \max_k \sigma_k = \max\left(\frac{1}{100}, \frac{1}{13}, \frac{1}{0.5}\right) = 2$$

$$\text{cond}_2(A) = \|A\|_2 \|A^{-1}\|_2 = 200$$

in general $\text{cond}_2(D) = \frac{D_{\max}}{D_{\min}}$

“Little c” demo

Discuss what happens when c is “close” to zero
What are the eigenvalues of triangular matrices?

only the first part

Remarks:

The need for pivoting does not depend on whether the matrix is singular.

A non-singular matrix always has a solution.

A singular matrix may not have a solution, or may have infinitely many solutions.

About condition numbers

1. For any matrix \mathbf{A} , $\text{cond}(\mathbf{A}) \geq 1$
2. For the identity matrix \mathbf{I} , $\text{cond}(\mathbf{I}) = 1$
3. For any matrix \mathbf{A} and a nonzero scalar γ , $\text{cond}(\gamma\mathbf{A}) = \text{cond}(\mathbf{A})$
4. For any diagonal matrix \mathbf{D} , $\text{cond}(\mathbf{D}) = \frac{\max|d_i|}{\min|d_i|}$
5. The condition number is a measure of how close a matrix is to being singular: a matrix with large condition number is nearly singular, whereas a matrix with a condition number close to 1 is far from being singular
6. The determinant of a matrix is NOT a good indicator if a matrix is near singularity
if $\det(\mathbf{A}) = 0 \rightarrow \text{singular}$

Residual versus error

Our goal is to find the solution \mathbf{x} to the linear system of equations $\mathbf{A} \mathbf{x} = \mathbf{b}$

Let us recall the solution of the perturbed problem

$$\hat{\mathbf{x}} = (\mathbf{x} + \Delta\mathbf{x})$$

approx → *exact*

which could be the solution of

$$\mathbf{A} \hat{\mathbf{x}} = (\mathbf{b} + \Delta\mathbf{b}), \quad (\mathbf{A} + \mathbf{E}) \hat{\mathbf{x}} = \mathbf{b}, \quad (\mathbf{A} + \mathbf{E}) \hat{\mathbf{x}} = (\mathbf{b} + \Delta\mathbf{b})$$

And the **error vector** as

$$\mathbf{e} = \Delta\mathbf{x} = \hat{\mathbf{x}} - \mathbf{x}$$

→ but we usually don't have \approx ! !

We can write the **residual vector** as

$$\mathbf{r} = \mathbf{b} - \mathbf{A} \hat{\mathbf{x}}$$

→ known quantity

BACK TO HILBERT

Residual versus error

Our goal is to find the solution \mathbf{x} to the linear system of equations $\mathbf{A} \mathbf{x} = \mathbf{b}$

It can be shown that the residual $\mathbf{r} = \mathbf{b} - \mathbf{A} \hat{\mathbf{x}}$ satisfies the following:

$$\frac{\|\mathbf{r}\|}{\|\mathbf{A}\| \|\hat{\mathbf{x}}\|} \leq c \epsilon_m$$

Where c is large without pivoting and small with partial pivoting.

Therefore, solving the system of equations with partial pivoting yields **small relative residual regardless of conditioning of the system.**

Residual versus error

Relative residual: $\frac{\|r\|}{\|A\|\|x\|}$ (How well the solution satisfies the problem)

This is what I can measure

Relative error: $\frac{\|\Delta x\|}{\|x\|}$ (How close the approximated solution is from the exact one)

This is what I am interested in!

When solving a system of linear equations via LU with partial pivoting, the relative residual is guaranteed to be small!

For well-conditioned matrices, small relative residual implies small relative error.

$$\frac{\|\Delta x\|}{\|x\|} \leq \text{cond}(A) \frac{\|r\|}{\|A\|\|x\|}$$

When solving $AX=b$ with LU + partial pivoting, the error will be small for well-conditioned A !

Rule of thumb for conditioning

Suppose we want to find the solution \mathbf{x} to the linear system of equations $\mathbf{A} \mathbf{x} = \mathbf{b}$ using LU factorization with partial pivoting and backward/forward substitutions.

Suppose we compute the solution $\hat{\mathbf{x}}$.

If the entries in \mathbf{A} and \mathbf{b} are accurate to S decimal digits,

and $\text{cond}(\mathbf{A}) = 10^W$,

then the elements of the solution vector $\hat{\mathbf{x}}$ will be accurate to about

$$S - W$$

decimal digits

Clicker question

DEMO

Matrix Conditioning: Accurate digits

1 point

Let's say we want to solve the following linear system:

$$Ax = b$$

Assuming you are working with IEEE double precision floating point numbers, how many digits of accuracy will your answer have if $\kappa(A) = 1000$?

- A) 3
- B) 10
- C) 13
- D) 16
- E) 32

A, b → 16 digits of accuracy

$$\text{cond}(\underline{A}) = 10^3$$

Hence \underline{x} will have $16 - 3 = 13$ digits of accuracy.