

Introduction to Simulation - Lecture 3

Basics of Solving Linear Systems

Jacob White

Thanks to Deepak Ramaswamy, Michal Rewienski,
Karen Veroy and Jacob White

SMA-HPC ©2003 MIT

Outline

Solution Existence and Uniqueness

Gaussian Elimination Basics

- LU factorization

- Pivoting and Growth

Hard to solve problems

- Conditioning

Application Problems

$$\underbrace{\begin{bmatrix} G \end{bmatrix}}_M \underbrace{V_n}_x = \underbrace{I_s}_b$$

- No voltage sources or rigid struts
- Symmetric and Diagonally Dominant
- Matrix is $n \times n$

Systems of Linear Equations

$$\begin{bmatrix} \uparrow & \uparrow & \dots & \uparrow \\ \vec{M}_1 & \vec{M}_2 & \dots & \vec{M}_N \\ \downarrow & \downarrow & \dots & \downarrow \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix}$$



$$x_1 \vec{M}_1 + x_2 \vec{M}_2 + \dots + x_N \vec{M}_N = b$$

Find a set of weights, x , so that the weighted sum of the columns of the matrix M is equal to the right hand side b

Systems of Linear Equations

Key Questions

- Given $Mx = b$
 - Is there a solution?
 - Is the solution Unique?
- Is there a Solution?

There exists weights, x_1, \dots, x_N , such that

$$x_1 \vec{M}_1 + x_2 \vec{M}_2 + \dots + x_N \vec{M}_N = b$$

A solution exists when b is in the span of the columns of M

- Is the Solution Unique?

Suppose there exists weights, y_1, \dots, y_N , not all zero

$$y_1 \vec{M}_1 + y_2 \vec{M}_2 + \dots + y_N \vec{M}_N = \mathbf{0}$$

Then if $Mx = b$, therefore $M(x + y) = b$

A solution is unique only if the columns of M are linearly independent.

- Given $Mx = b$, where M is square
 - If a solution exists for any b , then the solution for a specific b is unique.

For a solution to exist for any b , the columns of M must span all N -length vectors. Since there are only N columns of the matrix M to span this space, these vectors must be linearly independent.

A square matrix with linearly independent columns is said to be nonsingular.

Gaussian Elimination Basics

Important Properties

Gaussian Elimination Method for Solving $Mx = b$

- A “Direct” Method
 - Finite Termination for exact result (ignoring roundoff)
- Produces accurate results for a broad range of matrices
- Computationally Expensive

Gaussian Elimination Basics

Reminder by Example

3 x 3 example

$$\begin{bmatrix} M_{11} & M_{12} & M_{13} \\ M_{21} & M_{22} & M_{23} \\ M_{31} & M_{32} & M_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$



$$\begin{aligned} M_{11}x_1 + M_{12}x_2 + M_{13}x_3 &= b_1 \\ M_{21}x_1 + M_{22}x_2 + M_{23}x_3 &= b_2 \\ M_{31}x_1 + M_{32}x_2 + M_{33}x_3 &= b_3 \end{aligned}$$

Gaussian Elimination Basics

Reminder by Example

Key Idea

Use Eqn 1 to Eliminate x_1 From Eqn 2 and 3

$$M_{11}x_1 + M_{12}x_2 + M_{13}x_3 = b_1$$

$$\left(M_{22} - \frac{M_{21}}{M_{11}} M_{12} \right) x_2 + \left(M_{23} - \frac{M_{21}}{M_{11}} M_{13} \right) x_3 = b_2 - \frac{M_{21}}{M_{11}} b_1$$

$$\left(M_{32} - \frac{M_{31}}{M_{11}} M_{12} \right) x_2 + \left(M_{33} - \frac{M_{31}}{M_{11}} M_{13} \right) x_3 = b_3 - \frac{M_{31}}{M_{11}} b_1$$

Gaussian Elimination Basics

Reminder by Example

Key Idea in the Matrix

Pivot

MULTIPLIERS

$$\begin{bmatrix} M_{11} \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} M_{12} \\ \left(M_{22} - \left(\frac{M_{21}}{M_{11}} \right) M_{12} \right) \\ \left(M_{32} - \left(\frac{M_{31}}{M_{11}} \right) M_{12} \right) \end{bmatrix} \begin{bmatrix} M_{13} \\ \left(M_{23} - \left(\frac{M_{21}}{M_{11}} \right) M_{13} \right) \\ \left(M_{33} - \left(\frac{M_{31}}{M_{11}} \right) M_{13} \right) \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 - \left(\frac{M_{21}}{M_{11}} \right) b_1 \\ b_3 - \left(\frac{M_{31}}{M_{11}} \right) b_1 \end{bmatrix}$$

Gaussian Elimination Basics

Reminder by Example

Remove x2 from Eqn 3

$$\begin{bmatrix} M_{11} & M_{12} & M_{13} \\ 0 & \left(M_{22} - \frac{M_{21}}{M_{11}} M_{12} \right) & \left(M_{23} - \frac{M_{21}}{M_{11}} M_{13} \right) \\ 0 & 0 & \left(M_{33} - \frac{M_{31}}{M_{11}} M_{13} - \frac{M_{32} - \frac{M_{21}}{M_{11}} M_{12}}{M_{22} - \frac{M_{21}}{M_{11}} M_{12}} \left(M_{23} - \frac{M_{21}}{M_{11}} M_{13} \right) \right) \end{bmatrix}$$

Pivot (points to $M_{22} - \frac{M_{21}}{M_{11}} M_{12}$)
Multiplier (points to $\frac{M_{32} - \frac{M_{21}}{M_{11}} M_{12}}{M_{22} - \frac{M_{21}}{M_{11}} M_{12}}$)

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 - \frac{M_{21}}{M_{11}} b_1 \\ b_3 - \frac{M_{31}}{M_{11}} b_1 - \frac{M_{32} - \frac{M_{21}}{M_{11}} M_{12}}{M_{22} - \frac{M_{21}}{M_{11}} M_{12}} \left(b_2 - \frac{M_{21}}{M_{11}} b_1 \right) \end{bmatrix}$$

Gaussian Elimination Basics

Reminder by Example

Simplify the Notation

$$\begin{bmatrix} M_{11} & M_{12} & M_{13} \\ 0 & \tilde{M}_{22} & \tilde{M}_{23} \\ 0 & \tilde{M}_{32} & \tilde{M}_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ \tilde{b}_2 \\ \tilde{b}_3 \end{bmatrix}$$

Gaussian Elimination Basics

Reminder by Example

Remove x_2 from Eqn 3

$$\begin{bmatrix} M_{11} & M_{12} & M_{13} \\ 0 & \tilde{M}_{22} & \tilde{M}_{23} \\ 0 & 0 & \tilde{M}_{33} - \frac{\tilde{M}_{32}}{\tilde{M}_{22}} \tilde{M}_{23} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ \tilde{b}_2 \\ \tilde{b}_3 - \frac{\tilde{M}_{32}}{\tilde{M}_{22}} \tilde{b}_2 \end{bmatrix}$$

Pivot

Multiplier

SMA-HPC ©2003 MIT

Gaussian Elimination Basics

Reminder by Example

GE Yields Triangular System

$$\begin{bmatrix} U_{11} & U_{12} & U_{13} \\ 0 & U_{22} & U_{23} \\ 0 & 0 & U_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}$$

$$x_3 = \frac{y_3}{U_{33}}$$

$$x_2 = \frac{y_2 - U_{23}x_3}{U_{22}}$$

$$x_1 = \frac{y_1 - U_{12}x_2 - U_{13}x_3}{U_{11}}$$

Gaussian Elimination Basics

Reminder by Example

The right-hand side updates

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 - \left(\frac{M_{21}}{M_{11}}\right) b_1 \\ b_3 - \left(\frac{M_{31}}{M_{11}}\right) b_1 - \frac{\tilde{M}_{32}}{\tilde{M}_{22}} \tilde{b}_2 \end{bmatrix} \quad \rightarrow \quad \begin{bmatrix} 1 & 0 & 0 \\ \frac{M_{21}}{M_{11}} & 1 & 0 \\ \frac{M_{31}}{M_{11}} & \frac{\tilde{M}_{32}}{\tilde{M}_{22}} & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

Gaussian Elimination Basics

Reminder by Example

Fitting the pieces together

$$\begin{bmatrix} M_{11} & M_{12} & M_{13} \\ \frac{M_{21}}{M_{11}} & \tilde{M}_{22} & \tilde{M}_{23} \\ \frac{M_{31}}{M_{11}} & \frac{\tilde{M}_{32}}{\tilde{M}_{22}} & \tilde{M}_{33} \end{bmatrix}$$

Basics of LU Factorization

Solve $Mx = b$

Step 1

$$M = L \cdot U$$



Step 2 Forward Elimination

Solve $Ly = b$

Step 3 Backward Substitution

Solve $Ux = y$

SMA-HPC ©2003 MIT

Recall from basic linear algebra that a matrix A can be factored into the product of a lower and an upper triangular matrix using Gaussian Elimination. The basic idea of Gaussian elimination is to use equation one to eliminate x_1 from all but the first equation. Then equation two is used to eliminate x_2 from all but the second equation. This procedure continues, reducing the system to upper triangular form as well as modifying the right hand side. More is needed here on the basics of Gaussian elimination.

Basics of LU Factorization

Solving Triangular Systems

Matrix

$$\begin{bmatrix} l_{11} & 0 & \dots & 0 \\ l_{21} & l_{22} & \dots & 0 \\ l_{13} & \dots & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ l_{N1} & \dots & \dots & l_{NN} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix}$$

The First Equation has only y_1 as an unknown.
The Second Equation has only y_1 and y_2 as unknowns.

Basics of LU Factorization

Solving Triangular Systems

Algorithm

$$\begin{bmatrix} l_{11} & 0 & \dots & 0 \\ l_{21} & l_{22} & & \\ l_{13} & & & \\ \vdots & & & \\ l_{N1} & & & l_{NN} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix}$$

$$\begin{aligned}
 y_1 &= \frac{1}{l_{11}} b_1 \\
 y_2 &= \frac{1}{l_{22}} (b_2 - l_{21} y_1) \\
 y_3 &= \frac{1}{l_{33}} (b_3 - l_{31} y_1 - l_{32} y_2) \\
 &\vdots \\
 y_N &= \frac{1}{l_{NN}} (b_N - \sum_{i=1}^{N-1} l_{Ni} y_i)
 \end{aligned}$$

SMA-HPC ©2003 MIT

Solving a triangular system of equations is straightforward but expensive. y_1 can be computed with one divide, y_2 can be computed with a multiply, a subtraction and a divide. Once y_{k-1} has been computed, y_k can be computed with $k-1$ multiplies, $k-2$ adds, a subtraction and a divide. Roughly the number of arithmetic operations is

$$N \text{ divides} + 0 \text{ add/subs} + 1 \text{ add/subs} + \dots + N-1 \text{ add/sub}$$

$$\begin{matrix} \uparrow & \uparrow & \uparrow \\ \text{for } y_1 & \text{for } y_2 & \text{for } y_N \end{matrix}$$

$$+ 0 \text{ mults} + 1 \text{ mult} + \dots + N-1 \text{ mults}$$

$$\begin{matrix} \uparrow & \uparrow & \uparrow \\ \text{for } y_1 & \text{for } y_2 & \text{for } y_N \end{matrix}$$


$$= (N-1)(N-2)/2 \text{ add/subs} + (N-1)(N-2)/2 \text{ mults} + N \text{ divides}$$

$$= \text{Order } N^2 \text{ operations}$$

Basics of LU Factorization

Factoring

Picture



The above is an animation of LU factorization. In the first step, the first equation is used to eliminate x_1 from the 2nd through 4th equation. This involves multiplying row 1 by a multiplier and then subtracting the scaled row 1 from each of the target rows. Since such an operation would zero out the a_{21} , a_{31} and a_{41} entries, we can replace those zero'd entries with the scaling factors, also called the multipliers. For row 2, the scale factor is a_{21}/a_{11} because if one multiplies row 1 by a_{21}/a_{11} and then subtracts the result from row 2, the resulting a_{21} entry would be zero. Entries a_{22} , a_{23} and a_{24} would also be modified during the subtraction and this is noted by changing the color of these matrix entries to blue. As row 1 is used to zero a_{31} and a_{41} , a_{31} and a_{41} are replaced by multipliers. The remaining entries in rows 3 and 4 will be modified during this process, so they are recolored blue.

This factorization process continues with row 2. Multipliers are generated so that row 2 can be used to eliminate x_2 from rows 3 and 4, and these multipliers are stored in the zero'd locations. Note that as entries in rows 3 and 4 are modified during this process, they are converted to green. The final step is to use row 3 to eliminate x_3 from row 4, modifying row 4's entry, which is denoted by converting a_{44} to pink.

It is interesting to note that as the multipliers are standing in for zero'd matrix entries, they are not modified during the factorization.

LU Basics

Factoring

Algorithm

```
For i = 1 to n-1 {           "For each Row"  
  For j = i+1 to n {       "For each target Row below the source"  
     $M_{ji} = \frac{M_{ji}}{M_{ii}}$  Pivot  
    For k = i+1 to n {    "For each Row element beyond Pivot"  
       $M_{jk} \leftarrow M_{jk} - M_{ji} M_{ik}$   
    }  
  }  
}
```


SMA-HPC ©2003 MIT

LU Basics

Factoring

Zero Pivots

At Step i



What if $M_{ii} = 0$? Cannot form $\frac{M_{ji}}{M_{ii}}$

Simple Fix (Partial Pivoting)

if $M_{ii} = 0$

Find $M_{ji} \neq 0 \quad j > i$

Swap Row j with i

SMA-HPC ©2003 MIT

Swapping row j with row i at step i of LU factorization is identical to applying LU to a matrix with its rows swapped “a priori”.

To see this consider swapping rows before beginning LU .

$$\begin{array}{c}
 \left[\begin{array}{c} \text{---} r_1 \text{---} \\ \text{---} r_2 \text{---} \\ \text{---} r_j \text{---} \\ \text{---} r_i \text{---} \\ \text{---} r_N \text{---} \end{array} \right] \begin{bmatrix} x_1 \\ x_2 \\ | \\ x_i \\ | \\ x_j \\ | \\ x_N \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ | \\ b_j \\ | \\ b_i \\ | \\ b_N \end{bmatrix}
 \end{array}$$

Curved arrows on the left and right sides of the matrix equation indicate the swapping of rows j and i in both the coefficient matrix and the right-hand side vector.

Swapping rows corresponds to reordering only the equations. Notice that the vector of unknowns is NOT reordered.

Two Important Theorems


- 1) Partial pivoting (swapping rows) always succeeds if M is non singular
- 2) LU factorization applied to a strictly diagonally dominant matrix will never produce a zero pivot

SMA-HPC ©2003 MIT

Theorem Gaussian Elimination applied to strictly diagonally dominant matrices will never produce a zero pivot.

Proof

- 1) Show the first step succeeds.
- 2) Show the $(n-1) \times (n-1)$ sub matrix



is still strictly diagonally dominant.

First Step

Second row after first step $a_{11} \neq 0$ as $|a_{11}| > \sum_{j=2}^n |a_{1j}|$

$$0, a_{22} - \frac{a_{21}}{a_{11}} a_{12}, a_{23} - \frac{a_{21}}{a_{11}} a_{13}, \dots, a_{2n} - \frac{a_{21}}{a_{11}} a_{1n}$$

Is

$$\left| a_{22} - \frac{a_{21}}{a_{11}} a_{12} \right| > \left| a_{2j} - \frac{a_{21}}{a_{11}} a_{1j} \right| ?$$

Contrived Example

$$\begin{bmatrix} 10^{-17} & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

$$L = \begin{bmatrix} 1 & 0 \\ 10^{17} & 1 \end{bmatrix} \quad U = \begin{bmatrix} 10^{-17} & 1 \\ 0 & -\underbrace{10^{17} + 2}_{\text{small pivot}} \end{bmatrix}$$

Can we represent this ?

SMA-HPC ©2003 MIT

In order to compute the exact solution first forward eliminate

$$\begin{bmatrix} 1 & 0 \\ 10^{17} & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

and therefore $y_1 = 1$, $y_2 = 3 - 10^{17}$.

Backward substitution yields

$$\begin{bmatrix} 10^{-17} & 1 \\ 0 & 2 - 10^{17} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 - 10^{17} \end{bmatrix}$$

and therefore $x_2 = \frac{3 - 10^{17}}{2 - 10^{17}} \approx +1$

and $x_1 = 10^{17} \left(1 - \frac{3 - 10^{17}}{2 - 10^{17}} \right) = 10^{17} \left(\frac{2 - 10^{17} - (3 - 10^{17})}{2 - 10^{17}} \right) \approx +1$


In the rounded case

$$\begin{bmatrix} 1 & 0 \\ 10^{17} & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix} \quad y_1 = 1 \quad y_2 = -10^{17}$$

$$\begin{bmatrix} 10^{-17} & 1 \\ 0 & 10^{17} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ -10^{17} \end{bmatrix} \quad x_1 = 1 \quad x_2 = 0$$

An Aside on Floating Point Arithmetic

Double precision number $X.XXX...X \cdot 10^{\text{exponent}}$



Basic Problem

$$1.0 - 0.0000000000000001 = 1.0$$

or

$$1.0000001 - 1 + (\pi \cdot 10^{-7}) \text{ right 8 digits}$$

Key Issue

Avoid small differences between large numbers !

Back to the contrived example

$$LU_{\text{Exact}} = \begin{bmatrix} 1 & 0 \\ 10^{17} & 1 \end{bmatrix} \begin{bmatrix} 10^{-17} & 1 \\ 0 & 2 - 10^{17} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

$$LU_{\text{Rounded}} = \begin{bmatrix} 1 & 0 \\ 10^{17} & 1 \end{bmatrix} \begin{bmatrix} 10^{-17} & 1 \\ 0 & -10^{17} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}_{\text{Exact}} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}_{\text{Rounded}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

LU Basics

Numerical Problems

Small Pivots

Partial Pivoting for Roundoff reduction

$$\text{If } |M_{ii}| < \max_{j>i} |M_{ji}|$$

Swap row i with $\arg(\max_{j>i} |M_{ij}|)$

$$\text{LU}_{\text{reordered}} = \begin{bmatrix} 1 & 0 \\ 10^{-17} & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & 1 + 2 \cdot 10^{-17} \end{bmatrix}$$

This multiplier is small

This term gets rounded

SMA-HPC ©2003 MIT

To see why pivoting helped notice that

$$\begin{bmatrix} 1 & 0 \\ 10^{-17} & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$$

yields $y_1 = 3$, $y_2 = 1 - 10^{-17} \approx 1$

Notice that without partial pivoting y_2 was $3 \cdot 10^{-17}$ or $\approx -10^{-17}$ with rounding. The right hand side value 3 in the unpivoted case was rounded away, where as now it is preserved. Continuing with the back substitution.

$$\begin{bmatrix} 1 & 2 \\ 0 & 1 + 2 \cdot 10^{-17} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$$

$x_2 \approx 1$ $x_1 \approx 1$

If the matrix is diagonally dominant or partial pivoting for round-off reduction is during LU Factorization:

- 1) The multipliers will always be smaller than one in magnitude.
- 2) The maximum magnitude entry in the LU factors will never be larger than $2^{(n-1)}$ times the maximum magnitude entry in the original matrix.

SMA-HPC ©2003 MIT

To see why pivoting helped notice that

$$\begin{bmatrix} 1 & 0 \\ 10^{-17} & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$$

yields $y_1 = 3$, $y_2 = 1 - 10^{-17} \approx 1$

Notice that without partial pivoting y_2 was $3 \cdot 10^{-17}$ or $\approx -10^{-17}$ with rounding. The right hand side value 3 in the unpivoted case was rounded away, where as now it is preserved. Continuing with the back substitution.

$$\begin{bmatrix} 1 & 2 \\ 0 & 1 + 2 \cdot 10^{-17} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$$

$x_2 \approx 1$ $x_1 \approx 1$


Hard to Solve Systems

Fitting Example

Polynomial Interpolation

Table of Data

t_0	$f(t_0)$
t_1	$f(t_1)$
⋮	⋮
t_N	$f(t_N)$



Problem fit data with an N^{th} order polynomial

$$f(t) = \alpha_0 + \alpha_1 t + \alpha_2 t^2 + \dots + \alpha_N t^N$$

Hard to Solve Systems

Example Problem

Matrix Form


$$\underbrace{\begin{bmatrix} 1 & t_0 & t_0^2 & \dots & t_0^N \\ 1 & t_1 & t_1^2 & \dots & t_1^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & t_N & t_N^2 & \dots & t_N^N \end{bmatrix}}_{M_{\text{interp}}} \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_N \end{bmatrix} = \begin{bmatrix} f(t_0) \\ f(t_1) \\ \vdots \\ f(t_N) \end{bmatrix}$$

SMA-HPC ©2003 MIT

The k th row in the system of equations on the slide corresponds to insisting that the N th order polynomial match the data exactly at point t_k . Notice that we selected the order of the polynomial to match the number of data points so that a square system is generated. This would not generally be the best approach to fitting data, as we will see in the next slides.

Hard to Solve Systems

Fitting Example



SMA-HPC ©2003 MIT

Notice what happens when we try to fit a high order polynomial to a function that is nearly t . Instead of getting only one coefficient to be one and the rest zero, instead when 100th order polynomial is fit to the data, extremely large coefficients are generated for the higher order terms. This is due to the extreme sensitivity of the problem, as we shall see shortly.

Hard to Solve Systems

Perturbation Analysis

Measuring Error Norms

Vector Norms

L₂ (Euclidean) norm :




$$\|x\|_2 = \sqrt{\sum_{i=1}^n |x_i|^2}$$

L₁ norm :

$$\|x\|_1 = \sum_{i=1}^n |x_i|$$

L_∞ norm :

$$\|x\|_\infty = \max_i |x_i|$$



Hard to Solve Systems

Perturbation Analysis

Measuring Error Norms

Matrix Norms

Vector induced norm : $\|A\| = \max_x \frac{\|Ax\|}{\|x\|} = \max_{\|x\|=1} \|Ax\|$

Induced norm of A is the maximum "magnification" of x by A

Easy norms to compute:

$$\|A\|_1 = \max_j \sum_{i=1}^n |A_{ij}| = \text{max abs column sum}$$

Why? Let $x = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$

$$\|A\|_\infty = \max_i \sum_{j=1}^n |A_{ij}| = \text{max abs row sum}$$

Why? Let $x = \begin{bmatrix} \pm 1 \\ \vdots \\ \pm 1 \end{bmatrix}$

$$\|A\|_2 = \text{not so easy to compute!!}$$

Hard to Solve Systems

Perturbation Analysis

Perturbation Equation

$$\underbrace{(M + \delta M)}_{\substack{\text{Models LU} \\ \text{Roundoff}}} \underbrace{(x + \delta x)}_{\substack{\text{Models Solution} \\ \text{Perturbation}}} = Mx + \delta Mx + M\delta x + \delta M\delta x = \underbrace{b}_{\substack{\text{Unperturbed} \\ \text{RightHandSide}}}$$

Since $Mx - b = 0$

$$M\delta x = -\delta M(x + \delta x) \Rightarrow \delta x = -M^{-1}\delta M(x + \delta x)$$

Taking Norms $\|\delta x\| \leq \|M^{-1}\| \|M\| \frac{1}{\|M\|} \|\delta M\| \|x + \delta x\|$

Relative Error Relation $\frac{\|\delta x\|}{\|x + \delta x\|} \leq \underbrace{\|M^{-1}\| \|M\|}_{\substack{\text{"Condition} \\ \text{Number"}}} \frac{\|\delta M\|}{\|M\|}$

SMA-HPC ©2003 MIT

As the algebra on the slide shows the relative changes in the solution x is bounded by an M dependent factor times the relative changes in M . The factor

$$\|M^{-1}\| \|M\|$$


was historically referred to as the condition number of M , but that definition has been abandoned as then the condition number is norm dependent. Instead the condition number of M is the ratio of singular values of M .

$$\text{cond}(M) = \frac{\sigma_{\max}(M)}{\sigma_{\min}(M)}$$


Singular values are outside the scope of this course, consider consulting Trefethen & Bau.

Geometric Approach is clearer

$M = [\vec{M}_1 \ \vec{M}_2]$, Solving $Mx = b$ is finding $x_1 \vec{M}_1 + x_2 \vec{M}_2 = b$



Columns orthogonal




Columns nearly aligned

When vectors are nearly aligned, difficult to determine how much of \vec{M}_1 versus how much of \vec{M}_2


Hard to Solve Systems

Geometric Analysis

Polynomial Interpolation



The power series polynomials are nearly linearly dependent



SMA-HPC ©2003 MIT

Question Does row scaling reduce growth ?

$$\begin{bmatrix} d_{11} & 0 & \dots & 0 \\ 0 & \dots & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & 0 & d_{NN} \end{bmatrix} \begin{bmatrix} a_{11} & \dots & a_{1N} \\ \vdots & \vdots & \vdots \\ a_{N1} & \dots & a_{NN} \end{bmatrix} = \begin{bmatrix} d_{11} a_{11} & \dots & d_{11} a_{1N} \\ \vdots & \vdots & \vdots \\ d_{NN} a_{N1} & \dots & d_{NN} a_{NN} \end{bmatrix}$$

Does row scaling reduce condition number ?

$\| M \| \| M^{-1} \| \equiv$ condition number of M

Theorem If floating point arithmetic is used, then row scaling ($D M x = D b$) will not reduce growth in a meaningful way.

If

$$\hat{M}_{LU} \hat{x} = b$$

and

$$\widehat{D M}_{LU} \hat{x}' = D b$$

then

$$\hat{x} = \hat{x}' \leftarrow \text{No roundoff reduction}$$

Summary

Solution Existence and Uniqueness

Gaussian Elimination Basics

- LU factorization

- Pivoting and Growth

Hard to solve problems

- Conditioning