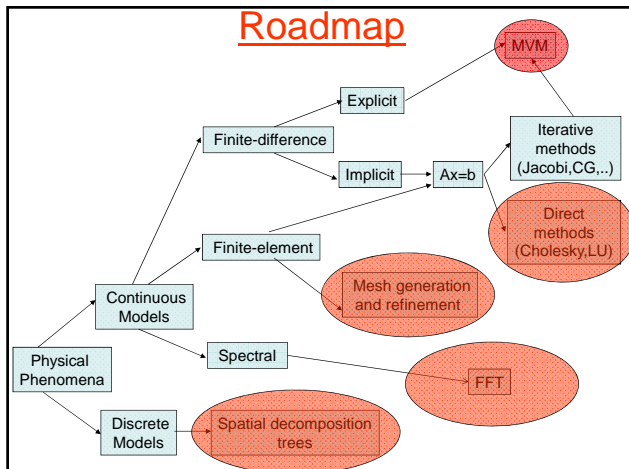


Some Computational Science Algorithms

Computational science

- Simulations of physical phenomena
 - fluid flow over aircraft (Boeing 777)
 - fatigue fracture in aircraft bodies
 - evolution of galaxies
 -
- Two main approaches
 - continuous models: fields and differential equations (eg. Navier-Stokes equations, Maxwell's equations,...)
 - discrete models: particles and forces (eg. gravitational forces)
- Paradox
 - most differential equations cannot be solved exactly
 - must use numerical techniques that convert calculus problem to matrix computations: **discretization**
 - n-body methods are straight-forward
 - but need to use a lot of bodies to get accuracy
 - must find a way to reduce $O(N^2)$ complexity of obvious algorithm

Roadmap




Organization

- Finite-difference methods
 - ordinary and partial differential equations
 - discretization techniques
 - explicit methods: Forward-Euler method
 - implicit methods: Backward-Euler method
- Finite-element methods
 - mesh generation and refinement
 - weighted residuals
- N-body methods
 - Barnes-Hut
- Key algorithms and data structures
 - matrix computations
 - algorithms
 - MVM and MMM
 - » solution of systems of linear equations
 - » direct methods
 - » iterative methods
 - data structures
 - dense and sparse matrices
 - graph computations
 - mesh generation and refinement
 - spatial decomposition trees

Ordinary differential equations

- Consider the ode
 $u'(t) = -3u(t)+2$
 $u(0) = 1$
- This is called an **initial value problem**
 - initial value of u is given
 - compute how function u evolves for $t > 0$
- Using elementary calculus, we can solve this ode exactly
 $u(t) = 1/3 (e^{-3t}+2)$




Problem

- For general ode's, we may not be able to express solution in terms of elementary functions
- In most practical situations, we do not need exact solution anyway
 - enough to compute an approximate solution, provided
 - we have some idea of how much error was introduced
 - we can improve the accuracy as needed
- General solution:
 - convert calculus problem into algebra/arithmetic problem
 - discretization: replace continuous variables with discrete variables
 - in finite differences,
 - time will advance in fixed-size steps: $t=0, h, 2h, 3h, \dots$
 - differential equation is replaced by difference equation

Forward-Euler method

- Intuition:
 - we can compute the derivative at $t=0$ from the differential equation
 $u'(t) = -3u(t)+2$
 - so compute the derivative at $t=0$ and advance along tangent to $t=h$ to find an approximation to $u(h)$
- Formally, we replace derivative with forward difference to get a difference equation
 - $u'(t) \rightarrow (u(t+h) - u(t))/h$
- Replacing derivative with difference is essentially the inverse of how derivatives were probably introduced to you in elementary calculus



Back to ode

- Original ode
 $u'(t) = -3u(t)+2$
- After discretization using Forward-Euler:
 $(u_i(t+h) - u_i(t))/h = -3u_i(t)+2$
- After rearrangement, we get **difference equation**
 $u_i(t+h) = (1-3h)u_i(t)+2h$
- We can now compute values of u :
 $u_i(0) = 1$
 $u_i(h) = (1-h)$
 $u_i(2h) = (1-2h+3h^2)$
 \dots

Exact solution of difference equation

- In this particular case, we can actually solve difference equation exactly

- It is not hard to show that if difference equation is

$$u_i(t+h) = a^* u_i(t) + b$$

$$u_i(0) = 1$$

the solution is

$$u_i(nh) = a^n + b(1-a^n)/(1-a)$$

- For our difference equation,


$$u_i(t+h) = (1-3h)u_i(t) + 2h$$

the exact solution is

$$u_i(nh) = 1/3(1-3h)^n + 2$$


Comparison

- Exact solution
 $u(t) = 1/3(e^{-3t} + 2)$
 $u(nh) = 1/3(e^{-3nh} + 2)$ (at time-steps)
- Forward-Euler solution
 $u_i(nh) = 1/3(1-3h)^n + 2$
- Use series expansion to compare
 $u(nh) = 1/3(1-3nh + 9/2 n^2 h^2 - \dots + 2)$
 $u_i(nh) = 1/3(1-3nh + n(n-1)/2 9h^2 + \dots + 2)$
 So error = $O(nh^2)$ (provided $h < 2/3$)
- Conclusion:
 - error per time step (local error) = $O(h^2)$
 - error at time $nh = O(nh^2)$
- In general, Forward-Euler converges only if time step is "small enough"




Choosing time step

- Time-step needs to be small enough to capture highest frequency phenomenon of interest
- Nyquist's criterion
 - sampling frequency must be at least twice highest frequency to prevent aliasing
 - for most finite-difference formulas, you need sampling frequencies (much) higher than the Nyquist criterion
- In practice, most functions of interest are not band-limited, so use
 - insight from application or
 - reduce time-step repeatedly till changes are not significant
- Fixed-size time-step can be inefficient if frequency varies widely over time interval
 - other methods like finite-elements permit variable time-steps as we will see later




Backward-Euler method

- Replace derivative with backward difference
 $u'(t) \rightarrow (u(t) - u(t-h))/h$
- For our ode, we get
 $u_b(t) - u_b(t-h)/h = -3u_b(t) + 2$
 which after rearrangement
 $u_b(t) = (2h + u_b(t-h))/(1+3h)$
- As before, this equation is simple enough that we can write down the exact solution:
 $u_b(nh) = ((1/(1+3h))^n + 2)/3$
- Using series expansion, we get
 $u_b(nh) = (1-3nh + (-n(n-1)/2) 9h^2 + \dots + 2)/3$
 $u_b(nh) = (1-3nh + 9/2 n^2 h^2 + 9/2 nh^2 + \dots + 2)/3$
 So error = $O(nh^2)$ (for any value of h)



Comparison

- Exact solution
 $u(t) = 1/3 (e^{-3t} + 2)$
 $u(nh) = 1/3 (e^{-3nh} + 2)$ (at time-steps)
- Forward-Euler solution
 $u_f(nh) = 1/3 (1 - 3h)^n + 2$
error = $O(nh^2)$ (provided $h < 2/3$)
- Backward-Euler solution
 $u_b(nh) = 1/3 ((1/(1+3h))^n + 2)$
error = $O(nh^2)$ (h can be any value you want)
- Many other discretization schemes have been studied in the literature
 - Runge-Kutta
 - Crank-Nicolson
 - Upwind differencing
 - ...




Red: exact solution
Blue: Backward-Euler solution (h=0.1)
Green: Forward-Euler solution (h=0.1)

Higher-order difference formulas

First derivatives:

- Forward-Euler: $y'(t) \rightarrow (y_f(t+h) - y_f(t)) / h$
- Backward-Euler: $y'(t) \rightarrow (y_b(t) - y_b(t-h)) / h$
- Centered: $y'(t) \rightarrow (y_c(t+h) - y_c(t-h)) / 2h$



Second derivatives:

- Forward: $y''(t) \rightarrow \frac{(y_f(t+2h) - y_f(t+h)) - (y_f(t+h) - y_f(t))}{h^2}$
 $= \frac{y_f(t+2h) - 2y_f(t+h) + y_f(t)}{h^2}$
- Backward: $y''(t) \rightarrow \frac{y_b(t) - 2y_b(t-h) + y_b(t-2h)}{h^2}$
- Centered: $y''(t) \rightarrow \frac{y_c(t+h) - 2y_c(t) + y_c(t-h)}{h^2}$

Systems of ode's

- Consider a system of coupled ode's of the form
 $u'(t) = a_{11}u(t) + a_{12}v(t) + a_{13}w(t) + c_1(t)$
 $v'(t) = a_{21}u(t) + a_{22}v(t) + a_{23}w(t) + c_2(t)$
 $w'(t) = a_{31}u(t) + a_{32}v(t) + a_{33}w(t) + c_3(t)$
- If we use Forward-Euler method to discretize this system, we get the following system of simultaneous equations
 $u_f(t+h) - u_f(t) / h = a_{11}u_f(t) + a_{12}v_f(t) + a_{13}w_f(t) + c_1(t)$
 $v_f(t+h) - v_f(t) / h = a_{21}u_f(t) + a_{22}v_f(t) + a_{23}w_f(t) + c_2(t)$
 $w_f(t+h) - w_f(t) / h = a_{31}u_f(t) + a_{32}v_f(t) + a_{33}w_f(t) + c_3(t)$

Forward-Euler (contd.)

Rearranging, we get

$$u_f(t+h) = (1 + ha_{11})u_f(t) + ha_{12}v_f(t) + ha_{13}w_f(t) + hc_1(t)$$

$$v_f(t+h) = ha_{21}u_f(t) + (1 + ha_{22})v_f(t) + ha_{23}w_f(t) + hc_2(t)$$

$$w_f(t+h) = ha_{31}u_f(t) + ha_{32}v_f(t) + (1 + ha_{33})w_f(t) + hc_3(t)$$

Introduce vector/matrix notation

$$\underline{u}(t) = [u(t) \ v(t) \ w(t)]^T$$

$$\underline{A} = \dots$$

$$\underline{c}(t) = [c_1(t) \ c_2(t) \ c_3(t)]^T$$

Vector notation

- Our systems of equations was

$$\begin{aligned} u_i(t+h) &= (1+ha_{11})^*u_i(t) + ha_{12}^*v_i(t) + ha_{13}^*w_i(t) + hc_1(t) \\ v_i(t+h) &= ha_{21}^*u_i(t) + (1+ha_{22})^*v_i(t) + ha_{23}^*w_i(t) + hc_2(t) \\ w_i(t+h) &= ha_{31}^*u_i(t) + ha_{32}^*v_i(t) + (1+ha_{33})^*w_i(t) + hc_3(t) \end{aligned}$$
- This system can be written compactly as follows

$$\underline{U}(t+h) = (I+hA)\underline{U}(t)+h\underline{C}(t)$$
- We can use this form to compute values of $\underline{U}(h), \underline{U}(2h), \underline{U}(3h), \dots$
- Forward-Euler is an example of **explicit method** of discretization
 - key operation: matrix-vector (MVM) multiplication
 - in principle, there is a lot of parallelism
 - $O(n^2)$ multiplications
 - $O(n)$ reductions
 - parallelism is independent of runtime values

Backward-Euler

- We can also use Backward-Euler method to discretize system of ode's

$$\begin{aligned} u_b(t)-u_b(t-h)/h &= a_{11}^*u_b(t) + a_{12}^*v_b(t) + a_{13}^*w_b(t) + c_1(t) \\ v_b(t)-v_b(t-h)/h &= a_{21}^*u_b(t) + a_{22}^*v_b(t) + a_{23}^*w_b(t) + c_2(t) \\ w_b(t)-w_b(t-h)/h &= a_{31}^*u_b(t) + a_{32}^*v_b(t) + a_{33}^*w_b(t) + c_3(t) \end{aligned}$$
- We can write this in matrix notation as follows

$$(I-hA)\underline{U}(t) = \underline{U}(t-h)+h\underline{C}(t)$$
- Backward-Euler is example of **implicit method** of discretization
 - key operation: solving a dense linear system $M\underline{x} = \underline{y}$
- How do we solve large systems of linear equations?
- Matrix $(I-hA)$ is often very sparse
 - Important to exploit sparsity in solving linear systems

Diversion: Solving linear systems

Solving linear systems

- Linear system: $A\underline{x} = \underline{b}$
- Two approaches
 - direct methods: Cholesky, LU with pivoting
 - factorize A into product of lower and upper triangular matrices $A = LU$
 - solve two triangular systems

$$\begin{aligned} L\underline{y} &= \underline{b} \\ U\underline{x} &= \underline{y} \end{aligned}$$
 - problems:
 - even if A is sparse, L and U can be quite dense ("fill")
 - no useful information is produced until the end of the procedure
 - iterative methods: Jacobi, Gauss-Seidel, CG, GMRES
 - guess an initial approximation \underline{x}_0 to solution
 - error is $A\underline{x}_0 - \underline{b}$ (called residual)
 - repeatedly compute better approximation \underline{x}_{k+1} from residual $(A\underline{x}_k - \underline{b})$
 - terminate when approximation is "good enough"

Iterative method: Jacobi iteration


- Linear system
 $4x+2y=8$
 $3x+4y=11$
- Exact solution is $(x=1, y=2)$
- Jacobi iteration for finding approximations to solution
 - guess an initial approximation
 - iterate
 - use first component of residual to refine value of x
 - use second component of residual to refine value of y
- For our example
 $x_{i+1} = x_i - (4x_i + 2y_i - 8)/4$
 $y_{i+1} = y_i - (3x_i + 4y_i - 11)/4$
 - for initial guess $(x_0, y_0=0)$

i	0	1	2	3	4	5	6	7
x	0	2	0.625	1.375	0.8594	1.1406	0.9473	1.0527
y	0	2.75	1.250	2.281	1.7188	2.1055	1.8945	2.0396

Jacobi iteration: general picture

- Linear system $Ax = b$
- Jacobi iteration
 $M^*x_{i+1} = (M-A)x_i + b$ (where M is the diagonal of A)
 This can be written as
 $x_{i+1} = x_i - M^{-1}(Ax_i - b)$
- Key operation:
 - matrix-vector multiplication
- Caveat:
 - Jacobi iteration does not always converge
 - even when it converges, it usually converges slowly
 - there are faster iterative methods available: CG, GMRES,...
 - what is important from our perspective is that key operation in all these iterative methods is **matrix-vector multiplication**

Sparse matrix representations




MVM with sparse matrices

Coordinate storage

for $P = 1$ to NZ do
 $Y(A.row(P)) = Y(A.row(P)) + A.val(P) * X(A.column(P))$

CRS storage


for $I = 1$ to N do
 for $JJ = A.rowptr(I)$ to $A.rowptr(I+1)-1$ do
 $Y(I) = Y(I) + A.val(JJ) * X(A.column(JJ))$




Finite-differences:pde's

Finite-difference methods for solving partial differential equations

- Basic ideas carry over unchanged
- Example: 2-d heat equation
 $\nabla^2 u / \cdot x^2 + \nabla^2 u / \cdot y^2 = f(x,y)$
 assume temperature at boundary is fixed
- Discretize domain using a regular NxN grid of pitch h
- Approximate derivatives as centered differences
 $\nabla^2 u / \cdot y^2 \rightarrow ((u(i,j+1) - u(i,j)) / h - (u(i,j) - u(i,j-1)) / h) / h$
 $\nabla^2 u / \cdot x^2 \rightarrow ((u(i+1,j) - u(i,j)) / h - (u(i,j) - u(i-1,j)) / h) / h$
- So we get a system of (N-1)x(N-1) difference equations in terms of the unknowns at the (N-1)x(N-1) interior points
 ; interior point (i,j)
 $u(i,j+1) + u(i,j-1) + u(i+1,j) + u(i-1,j) - 4u(i,j) = h^2 f(i,j,h)$
- This system can be solved using any of our methods.



Example




$$\nabla^2 u / \cdot x^2 + \nabla^2 u / \cdot y^2 = f(x,y)$$

Assume $f(x,y) = 0$

- Unknown temperatures are T1, T2, T3, T4
- Discretized equation at point 1:

$$\frac{T_2 - T_1}{h} - \frac{T_1 - 50}{h} + \frac{100 - T_1}{h} - \frac{T_1 - T_3}{h} = 0$$

Example (contd)



$$\nabla^2 u / \cdot x^2 + \nabla^2 u / \cdot y^2 = f(x,y)$$

Assume $f(x,y) = 0$

$$-4T_1 + T_2 + T_3 = -150$$

$$T_1 - 4T_2 + T_4 = -300$$

$$T_1 - 4T_3 + T_4 = -350$$

$$T_2 + T_3 - 4T_4 = -500$$


$$\begin{bmatrix} -4 & 1 & 1 & 0 \\ 1 & -4 & 0 & 1 \\ 1 & 0 & -4 & 1 \\ 0 & 1 & 1 & -4 \end{bmatrix} \begin{bmatrix} T_1 \\ T_2 \\ T_3 \\ T_4 \end{bmatrix} = \begin{bmatrix} -150 \\ -300 \\ -350 \\ -500 \end{bmatrix}$$

$$\text{Solution: } T_1 = 119, T_2 = 156, T_3 = 169, T_4 = 206$$

We could use an iterative method like Jacobi to solve such systems.
 Naïve approach: represent the sparse matrix in some sparse format like coordinate storage, and use a sparse MVM routine.

However, we can do better than this because matrix is known statically.

Example (contd.)



$$\begin{aligned} -4T_1 + T_2 + T_3 &= -150 \\ T_1 - 4T_2 + T_4 &= -300 \\ T_1 - 4T_3 + T_4 &= -350 \\ T_2 + T_3 - 4T_4 &= -500 \end{aligned}$$

Jacobi


$$-\frac{5}{4}u_x^2 + -\frac{5}{4}u_y^2 = f(x,y)$$

Assume $f(x,y) = 0$

$$\begin{aligned} -4T_1^{n+1} + T_2^n + T_3^n + 0 &= -150 \\ T_1^n - 4T_2^{n+1} + 0 + T_4^n &= -300 \\ T_1^n + 0 - 4T_3^{n+1} + T_4^n &= -350 \\ 0 + T_2^n + T_3^n - 4T_4^{n+1} &= -500 \end{aligned}$$

$$\begin{aligned} T_1^{n+1} &= \frac{1}{4}(T_2^n + T_3^n + 100 + 50) \\ T_2^{n+1} &= \frac{1}{4}(T_1^n + T_4^n + 100 + 200) \\ T_3^{n+1} &= \frac{1}{4}(T_1^n + T_4^n + 300 + 50) \\ T_4^{n+1} &= \frac{1}{4}(T_2^n + T_3^n + 300 + 200) \end{aligned}$$

- Use two dense arrays to store current and next T values
- No need for sparse matrices: matrix inlined into code



Solving partial differential equations contd.)

- System of $(N-1) \times (N-1)$ difference equations in terms of the unknowns at the $(N-1) \times (N-1)$ interior points

• interior point (i,j)
 $u(i,j+1) + u(i,j-1) + u(i+1,j) + u(i-1,j) - 4u(i,j) = h^2 f(ih,jh)$


- Matrix notation: use row-major (natural) order for u's

$$\begin{bmatrix} \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots \\ 0 \dots 1 \dots 0 \dots 1 \dots 0 \dots 0 \dots 0 \dots 1 \dots 0 \dots 1 \dots 0 \dots 0 \dots 1 \dots 0 \dots 1 \dots 0 \dots 0 \dots 1 \dots 0 \dots 1 \dots 0 \dots \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix} \begin{bmatrix} u(i-1,j) \\ \dots \\ u(i,j-1) \\ u(i,j) \\ u(i,j+1) \\ \dots \\ u(i+1,j) \end{bmatrix} = h^2 \begin{bmatrix} f(ih,jh) \\ \dots \\ f(ih,jh) \\ \dots \\ f(ih,jh) \end{bmatrix}$$

Pentadiagonal sparse matrix

Can be represented using specialized sparse matrix formats

Since matrix is sparse, we should use an iterative method like Jacobi.



Useful to change data structures

- Data structure:
 - pentadiagonal matrix can be inlined into code
 - values of u at a given time-step can be stored in a 2D array
 - use two arrays and use them for odd and even time-steps

- Algorithm:


for each interior point
 $u_{n+1}[i,j] = (u_n[i,j+1] + u_n[i,j-1] + u_n[i+1,j] + u_n[i-1,j]) - h^2 f(ih,jh) / 4$

- Known as stencil codes

- Example shown is Jacobi iteration with five-point stencil

- Parallelism

- all interior points can be computed in parallel
- parallelism is independent of runtime values




5-point stencil

Observations

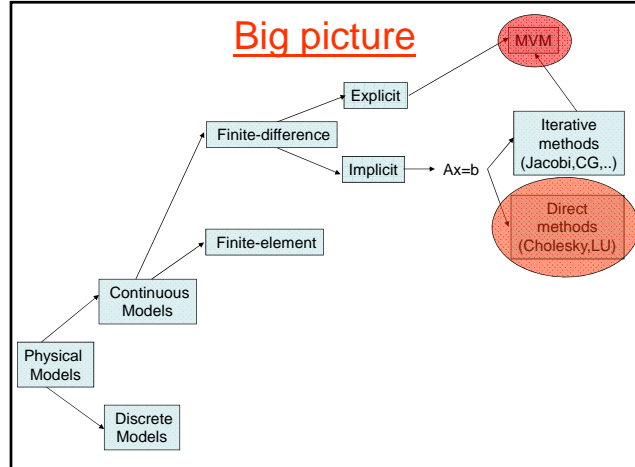
- Algorithm: Jacobi iteration with 5-point stencil to solve 2-D heat equation
- Two very different programs
 - a. pentadiagonal matrix: stored in sparse matrix format, unknowns: 1D vector
 - b. pentadiagonal matrix: inlined into code, unknowns: matrix
- Data structures are critical
 - can result in very different programs (implementations) for the same algorithm

Summary

- Finite-difference methods
 - can be used to find approximate solutions to ode's and pde's
- Many large-scale computational science simulations use these methods
- Time step or grid step needs to be constant and is determined by highest-frequency phenomenon
 - can be inefficient for when frequency varies widely in domain of interest
 - one solution: structured AMR methods




Big picture



Finite-element methods

- Express approximate solution to pde as a linear combination of certain basis functions
- Similar in spirit to Fourier analysis
 - express periodic functions as linear combinations of sines and cosines
- Questions:
 - what should be the basis functions?
 - mesh generation: discretization step for finite-elements
 - mesh defines basis functions ϕ_1, ϕ_2, \dots which are low-degree piecewise polynomial functions
 - given the basis functions, how do we find the best linear combination of these for approximating solution to pde?
 - $u = \sum c_i \phi_i$
 - weighted residual method: similar in spirit to what we do in Fourier analysis, but more complex because basis functions are not necessarily orthogonal

Mesh generation and refinement




- 1-D example:
 - mesh is a set of points, not necessarily equally spaced
 - basis functions are "hats" which
 - have a value of 1 at a mesh point,
 - decay down to 0 at neighboring mesh points
 - 0 everywhere else
 - linear combinations of these produce piecewise linear functions in domain, which may change slope only at mesh points
- In 2-D, mesh is a triangularization of domain, while in 3-D, it might be a tetrahedralization
- Mesh refinement: called h-refinement
 - add more points to mesh in regions where discretization error is large
 - irregular nature of mesh makes this easy to do this locally
 - finite-differences require global refinement which can be computationally expensive

Delaunay Mesh Refinement

- Iterative refinement to remove *bad* triangles with lots of discretization error:

```
while there are bad triangles do {
    Pick a bad triangle;
    Find its cavity;
    Retriangulate cavity;
    // may create new bad triangles
}
```

- Don't-care non-determinism:**
 - final mesh depends on order in which bad triangles are processed
 - applications do not care which mesh is produced
- Data structure:**
 - graph in which nodes represent triangles and edges represent triangle adjacencies
- Parallelism:**
 - bad triangles with cavities that do not overlap can be processed in parallel
 - parallelism is dependent on runtime values
 - compilers cannot find this parallelism
 - (Miller et al) at runtime, repeatedly build interference graph and find maximal independent sets for parallel execution



Finding coefficients


- Weighted residual technique**
 - similar in spirit to what we do in Fourier analysis, but basis functions are not necessarily orthogonal
- Key idea:**
 - problem is reduced to solving a system of equations $A\mathbf{x} = \mathbf{b}$
 - solution gives the coefficients in the weighted sum
 - because basis functions are zero almost everywhere in the domain, matrix A is usually very sparse
 - number of rows/columns of A ~ O(number of points in mesh)
 - number of non-zeros per row ~ O(connectivity of mesh point)
 - typical numbers:
 - A is $10^6 \times 10^6$
 - only about ~100 non-zeros per row

Finding the best choices of the coefficients:

Analogy with Fourier series:

$$f(x) = a_0 + \sum_i a_i \cos(ix) + \sum_i b_i \sin(ix)$$

How do you find 'best' choices for a's and b's?



$$\begin{aligned} \int_{-\pi}^{+\pi} f(x) \cos(kx) dx &= \int_{-\pi}^{+\pi} (a_0 + \sum_i a_i \cos(ix) + \sum_i b_i \sin(ix)) \cos(kx) dx \\ &= \int_{-\pi}^{+\pi} a_k \cos(kx) \cos(kx) dx \\ &= a_k \pi \end{aligned}$$

Key idea: - residual $f(x) - a_0 + \sum_i a_i \cos(ix) + \sum_i b_i \sin(ix)$
 - weight residual by known function and integrate to find corresponding coefficient

Weighted Residual Technique:

Residual: $(L u^* - f) = (L (\sum_i c_i \phi_i) - f)$

Weighted Residual $= (L (\sum_i c_i \phi_i) - f) \cdot \phi_k$

Equation for kth unknown: $\int_{\Omega} \phi_k * (L (\sum_i c_i \phi_i) - f) dV = 0 \Rightarrow$

If the differential equation is linear:


$$c_1 \int_{\Omega} \phi_k * L \phi_1 dV + \dots + c_N \int_{\Omega} \phi_k * L \phi_N dV = \int_{\Omega} \phi_k f dV \quad k = 1, 2, \dots, N$$

This system can be written as

$$Kc = b \text{ where } K(i,j) = \int_{\Omega} \phi_i * L \phi_j dV \quad b(i) = \int_{\Omega} \phi_i f dV$$

Key insight: Calculus problem of solving pde is converted to linear algebra problem of solving $Kc = b$ where K is sparse

Flow-chart of Adaptive Finite-element Simulation of Fracture



Barnes Hut N-body Simulation

Introduction

- Physical system simulation (time evolution)
 - System consists of **bodies**
 - “**n**” is the number of bodies
 - Bodies interact via **pair-wise forces**
- Many systems can be modeled in these terms
 - Galaxy clusters (gravitational force)
 - Particles (electric force, magnetic force)

Barnes Hut N-body Simulation

43

Barnes Hut Idea

- Precise force calculation
 - Requires $O(n^2)$ operations ($O(n^2)$ body pairs)
- Barnes and Hut (1986)
 - Algorithm to approximately compute forces
 - Bodies' initial position & velocity are also approximate
 - Requires only $O(n \log n)$ operations
 - Idea is to “combine” far away bodies
 - Error should be small because $force \sim 1/r^2$

Barnes Hut N-body Simulation

44


Barnes Hut Algorithm

- Set bodies' initial position and velocity
- Iterate over time steps
 1. Subdivide space until at most one body per cell
 - Record this spatial hierarchy in an octree
 2. Compute mass and center of mass of each cell
 3. Compute force on bodies by traversing octree
 - Stop traversal path when encountering a leaf (body) or an internal node (cell) that is far enough away
 4. Update each body's position and velocity

Barnes Hut N-body Simulation

45

Build Tree (Level 1)




Subdivide space until at most one body per cell

Barnes Hut N-body Simulation

46

Build Tree (Level 2)




Subdivide space until at most one body per cell

Barnes Hut N-body Simulation

47

Build Tree (Level 3)




Subdivide space until at most one body per cell

Barnes Hut N-body Simulation

48

Build Tree (Level 4)




Subdivide space until at most one body per cell

Barnes Hut N-body Simulation

49

Build Tree (Level 5)




Subdivide space until at most one body per cell

Barnes Hut N-body Simulation

50

Compute Cells' Center of Mass




For each internal cell, compute sum of mass and weighted average of position of all bodies in subtree; example shows two cells only

Barnes Hut N-body Simulation

51

Compute Forces




Compute force, for example, acting upon green body

Barnes Hut N-body Simulation

52

Compute Force (short distance)




Scan tree depth first from left to right; green portion already completed

Barnes Hut N-body Simulation

53

Compute Force (down one level)




Red center of mass is too close, need to go down one level

Barnes Hut N-body Simulation

54

Compute Force (long distance)




Yellow center of mass is far enough away

Barnes Hut N-body Simulation

55

Compute Force (skip subtree)



Therefore, entire subtree rooted in the yellow cell can be skipped

Barnes Hut N-body Simulation

56

Pseudocode

```

Set bodySet = ...
foreach timestep do {
  Octree octree = new Octree();
  foreach Body b in bodySet {
    octree.Insert(b);
  }
  OrderedList cellList = octree.CellsByLevel();
  foreach Cell c in cellList {
    c.Summarize();
  }
  foreach Body b in bodySet {
    b.ComputeForce(octree);
  }
  foreach Body b in bodySet {
    b.Advance();
  }
}

```

Barnes Hut N-body Simulation

57

Complexity

```

Set bodySet = ...
foreach timestep do {
  Octree octree = new Octree();           //  $O(n \log n)$ 
  foreach Body b in bodySet {             //  $O(n \log n)$ 
    octree.Insert(b);
  }
  OrderedList cellList = octree.CellsByLevel();
  foreach Cell c in cellList {           //  $O(n)$ 
    c.Summarize();
  }
  foreach Body b in bodySet {           //  $O(n \log n)$ 
    b.ComputeForce(octree);
  }
  foreach Body b in bodySet {           //  $O(n)$ 
    b.Advance();
  }
}

```

Barnes Hut N-body Simulation

58

Parallelism

```

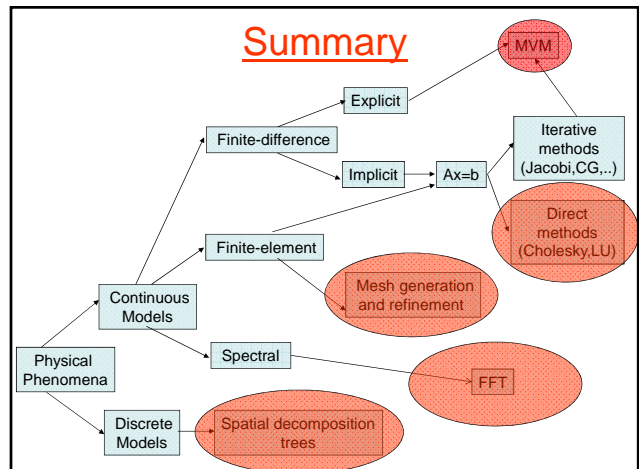
Set bodySet = ...
foreach timestep do {
  Octree octree = new Octree();           // sequential
  foreach Body b in bodySet {             // tree building
    octree.Insert(b);
  }
  OrderedList cellList = octree.CellsByLevel();
  foreach Cell c in cellList {           // tree traversal
    c.Summarize();
  }
  foreach Body b in bodySet {           // fully parallel
    b.ComputeForce(octree);
  }
  foreach Body b in bodySet {           // fully parallel
    b.Advance();
  }
}

```

Barnes Hut N-body Simulation

59

Summary



Summary (contd.)

- Some key computational science algorithms and data structures
 - MVM:
 - explicit finite-difference methods for ode's, iterative linear solvers, finite-element methods
 - both dense and sparse matrices
 - stencil computations:
 - finite-difference methods for pde's
 - dense matrices
 - $A=LU$:
 - direct methods for solving linear systems: factorization
 - usually only dense matrices
 - high-performance factorization codes use MMM as a kernel
 - mesh generation and refinement
 - finite-element methods
 - graphs

Summary (contd.)

- Terminology
 - regular algorithms:
 - dense matrix computations like MVM, $A=LU$, stencil computations
 - parallelism in algorithms is independent of runtime values, so all parallelization decisions can be made at compile-time
 - irregular algorithms:
 - graph computations like mesh generation and refinement
 - parallelism in algorithms is dependent on runtime values
 - most parallelization decisions have to be made at runtime during the execution of the algorithm
 - semi-regular algorithms:
 - sparse matrix computations like MVM, $A=LU$
 - parallelization decisions can be made at runtime once matrix is available, but before computation is actually performed
 - inspector-executor approach