CS601: Software Development for Scientific Computing

Autumn 2022

Week14 : Some important problems in motifs of scientific computing

(Minimum weight triangulation, Fast Multipole Method, Barnes-Hut)

Course Progress..

- **Pic source: the Parallel Computing Laboratory at U.C. Berkeley: A Research Agenda Based on the Berkeley View (2008)**

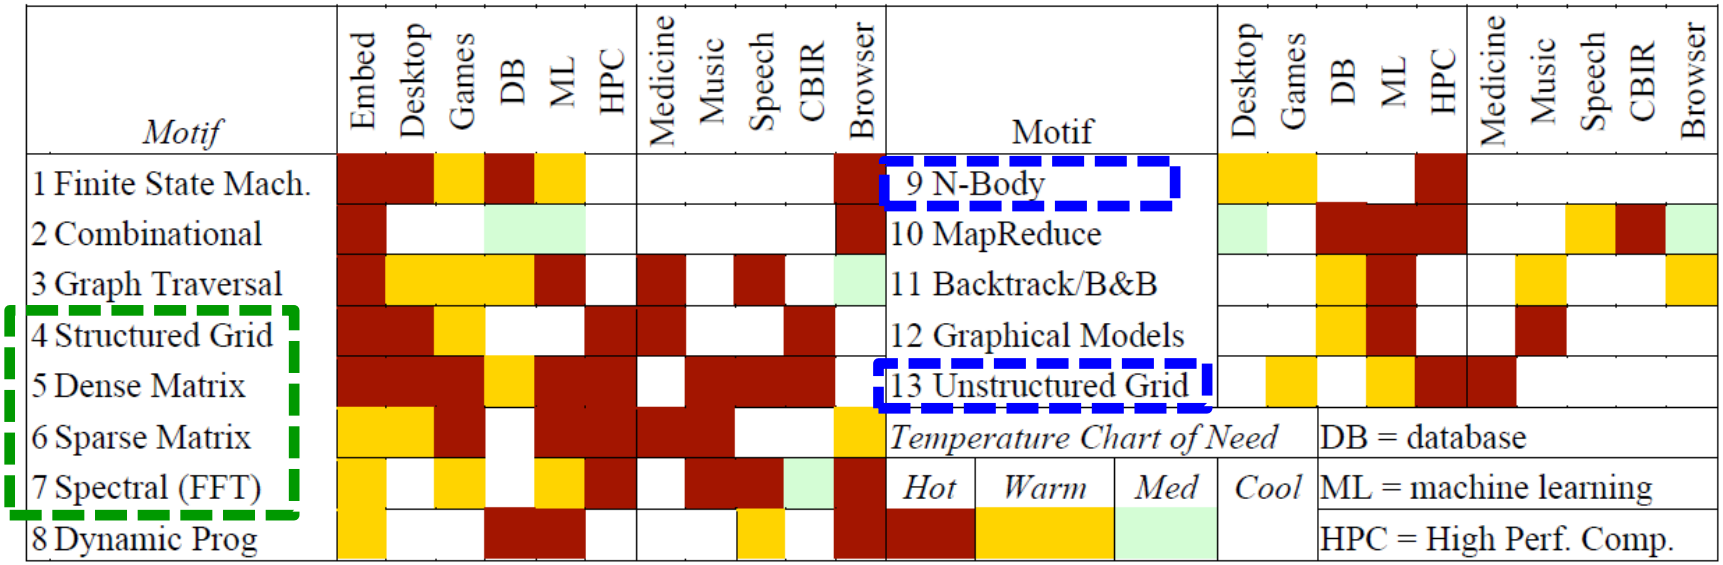

Figure 4. Temperature Chart of the 13 Motifs. It shows their importance to each of the original six application areas and then how important each one is to the five compelling applications of Section 3.1. More details on the motifs can be found in (Asanovic, Bodik et al. 2006).

Seen earlier

Next..

Unstructured Grids

- Motivation:
  - E.g. reduce the noise because of air flowing through a duct, Maintain uniformity in flow (no pressure drop)

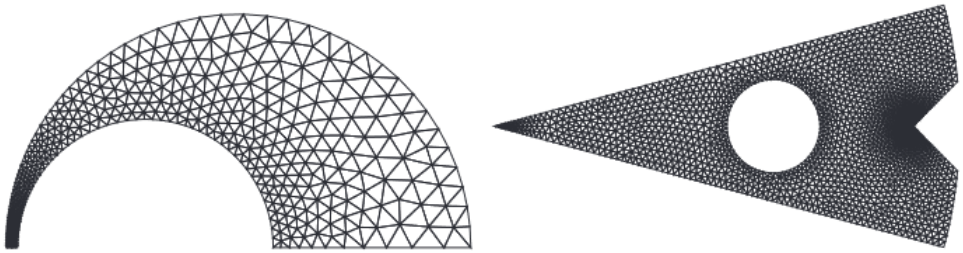

(e) Geometric Adaptivity (f) Pie with hole

https://www.math.uci.edu/~chenlong/Papers/Chen.L%3BHolst.M2010.pdf

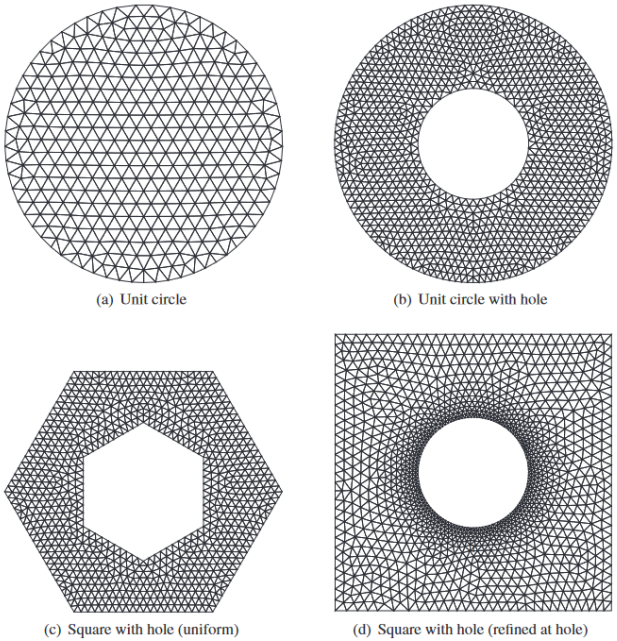

(a) Unit circle (b) Unit circle with hole

(c) Square with hole (uniform) (d) Square with hole (refined at hole)

- Handle complex geometries
- Refine at region of interest

Unstructured Grids

- Discretize the domain into an optimal set of triangles (or tetrahedra / simplex)
- Areas: Mesh Generation and Mesh Optimization

Iterate

1. Fix the location of vertices and optimize **OR**
2. Fix the vertices' connectivity and optimize (determine optimal placement of vertices)

(Optimize based on:

- Sum of edge weights

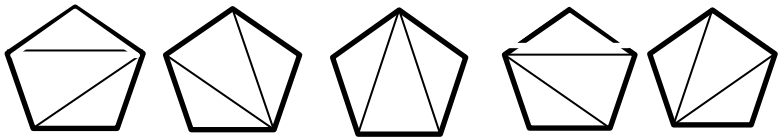

- Harmonic energy, Distortion energy)

Unstructured Grids – Program Representation (Examples)

- Option 1:

```
Point points[n]; // represents coordinates
int triangles[m][3]; // represents triangles
```

- Option 2:

```
double x[n], y[n]; // represents coordinates (2D)
int triangles[m][3]; // represents triangles
```

- Option 3: $G(V, E)$

```
double x[n], y[n]; // represents vertices/nodes (2D)
int edges[m][2]; // represents edges
```

Unstructured Grids - Challenges

- Placing the points
  - What constitutes inside and what is outside the domain? *Place points only in the interior.*
  - What should be the distance between points?
  - How many points should be there?
- Connecting the points
  - What is the best way to create tiles once the points are placed?

Unstructured Grids - Approaches

- Delaunay Triangulation is the most commonly used unstructured triangulation method
  - Advantage: can automatically give a ‘better’ triangulation (e.g. w.r.t aspect ratio)
  - Disadvantage: suitable for convex domain
- Advancing Front Method is another method
  - Advantage: suitable for concave domain
  - Disadvantage: No prioritization of triangulation

Minimum Weight Triangulation

- Type of divide-and-conquer with two properties:
  - Optimal substructure and repeated sub-problems.

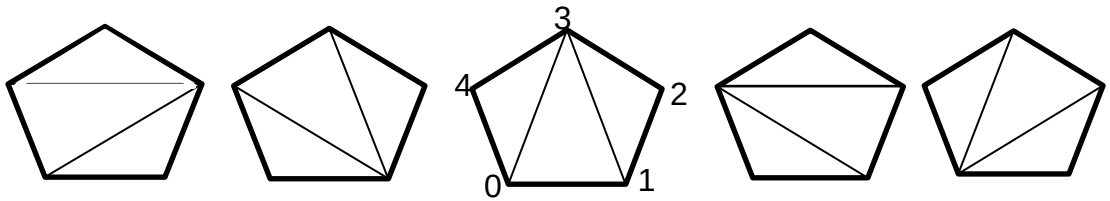

Minimum Weight Triangulation Problem

Objective: Triangulate a polygon such that edges do not intersect AND sum of edge lengths is minimized

$$C(i,j) = \begin{cases} min\left(C(i,j), \min\limits_{i<k<j} C(i,k) + C(k,j) + W(i,k,j)\right) \\ 0 & j \leq i+1 \\ Given\ W(i,j,k) \end{cases}$$

Minimum Weight Triangulation

- Pseudocode and call tree of triangulating a pentagon (vertices named 0 to 4)

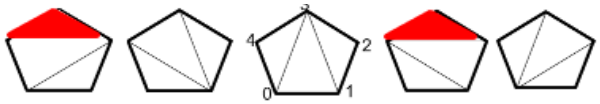

Minimum Weight Triangulation Problem

Objective: Triangulate a polygon such that edges do not intersect AND sum of edge lengths is minimized

$$C(i,j) = \begin{cases} min(\ C(i,j), \min\limits_{i<k<j} C(i,k) + C(k,j) + W(i,k,j)) \\ 0 & j \le i+1 \\ Given\ W(i,j,k) \end{cases}$$

```
Cost(i, j){
   if j < i+2 then return 0;
   curMin ← INFINITY;
   for k← i+1 to j-1 do
      res ← Cost(i,k) + Cost(k,j) + Weight(i,k,j)
      if res < curMin then curMin ← res;
   return curMin;
}
Weight(i,j,k){
   // return sum of Euclidian distances between
   // (i,j), (j,k), and (k,i)
}
```

```
main(){
   Cost(0,4);
}
```

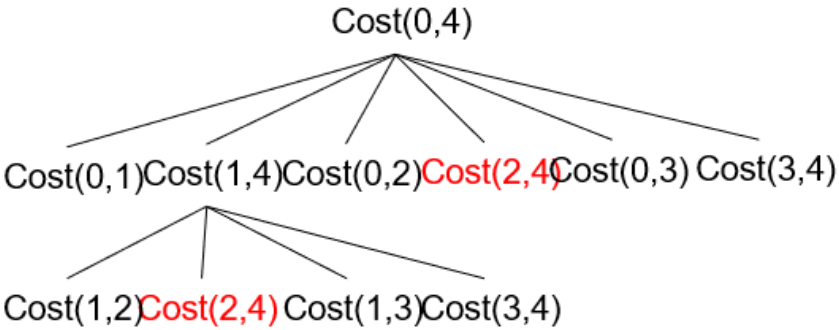

Minimum Weight Triangulation

- Iterative formulation (note the 2D array representing the matrix to be computed)

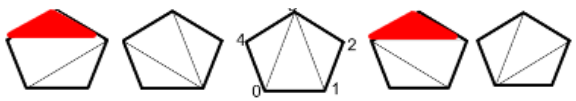

Minimum Weight Triangulation Problem

Objective: Triangulate a polygon such that edges do not intersect AND sum of edge lengths is minimized

$$C(i,j) = \begin{cases} \min\left(C(i,j), \min\limits_{i<k<j} C(i,k) + C(k,j) + W(i,k,j)\right) \\ 0 & j \le i+1 \\ Given\ W(i,j,k) \end{cases}$$

```
Cost(i, j){
   If j < i+2 then return 0;
   curMin ← INFINITY;
   for k← i+1 to j-1 do
      res ← Cost(i,k) + Cost(k,j) + Weight(i,k,j)
      If res < curMin then curMin ← res;
   return curMin;
}
Weight(i,j,k){
   // return sum of Euclidian distances between
   // (i,j), (j,k), and (k,i)
}
```

```
main(){
   Cost(0,4);
}
```

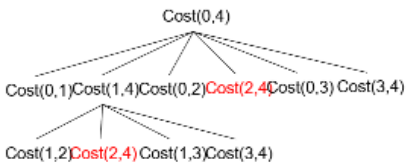


```
Cost(n){
  table[n][n]; //n is number of vertices
  for g ← 1 to n-1 do
     for i ← 0 to n-g do
        j ← i + g;
        table[i][j] ← INFINITY;
        for k ← i+1 to j–1 do
           res ← table[i][k] + table[k][j] + Weight(i,k,j)
           if res < table[i][j] then
              table[i][j] ← res;
}
```

Iterative formulation

Minimum Weight Triangulation

- The 2D array is used to compute only the upper triangular matrix. Cost of polygon (0,1,2,3) is shown.

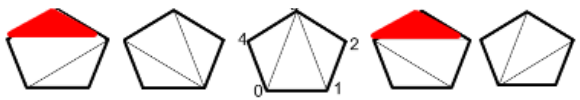

Minimum Weight Triangulation Problem

Objective: Triangulate a polygon such that edges do not intersect AND sum of edge lengths is minimized

$$C(i,j) = \begin{cases} \min\left(C(i,j), \min\limits_{i<k<j} C(i,k) + C(k,j) + W(i,k,j)\right) \\ 0 & j \le i+1 \\ Given\ W(i,j,k) \end{cases}$$

```
Cost(i, j){
   If j < i+2 then return 0;
   curMin ← INFINITY;
   for k← i+1 to j-1 do
      res ← Cost(i,k) + Cost(k,j) + Weight(i,k,j)
      If res < curMin then curMin ← res;
   return curMin;
}
Weight(i,j,k){
   // return sum of Euclidian distances between
   // (i,j), (j,k), and (k,i)
}
```

```
maIn(){
   Cost(0,4);
}
```

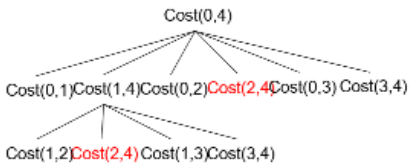


```
Cost(n){
  table[n][n]; //n is number of vertices
  for g ← 1 to n-1 do
     for i ← 0 to n-g do
        j ← i + g;
        table[i][j] ← INFINITY;
        for k ← i+1 to j–1 do
           res ← table[i][k] + table[k][j] + Weight(i,k,j)
           if res < table[i][j] then
              table[i][j] ← res;
}
```

Iterative formulation

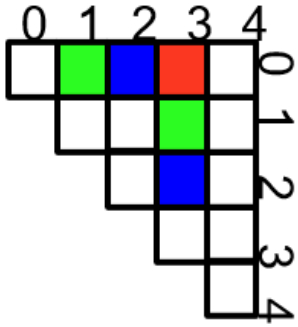

Cost(0,3) = min(Cost(0,1)+Cost(1,3)+ $\triangle^{le}$ (0,1,3),
Cost(0,2) + Cost(2,3) + $\triangle^{le}$ (0,2,3))

Further reading/viewing

- https://www.youtube.com/watch?v=lPcBX4BBW9U
- https://www.youtube.com/watch?v=tWf1z9i-Org
- https://www.math.uci.edu/~chenlong/Papers/Chen.L%3BHolst.M2010.pdf

Particle (Simulation) Methods

- N-Body Simulation – Problem

System of N-bodies (e.g. galaxies, stars, atoms, light rays etc.) interacting with each other continuously

  - Problem:
    - Compute force acting on a body due to all other bodies in the system
    - Determine position, velocity, at various times for each body
  - Objective:
    - Determine the (approximate) evolution of a system of bodies interacting with each other simultaneously

Particle (Simulation) Methods

- N-Body Simulation - Examples
  - Astrophysical simulation: E.g. each body is a star/galaxy https://commons.wikimedia.org/w/index.php?title=File%3AGalaxy_collision.ogv
  - Graphics: E.g. each body is a ray of light emanating from the light source. https://www.fxguide.com/fxfeatured/brave-new-hair/

    - Here each body is a point on a strand of hair

N-Body Simulation

- All-pairs Method
  - Naïve approach. Compute *all pair-wise interactions*
- Hierarchical Methods
  - Optimize. Reduce the number of pair-wise force calculations. How? dependence on ‘distant’ particle(s) can be *compressed*
  - Examples:
    - Barnes-Hut
    - Fast Multipole Method

N-Body Simulation

- Three fundamental simulation approaches
  - Particle-Particle (PP)
  - Particle-Mesh (PM)
  - Particle-Particle-Particle-Mesh (P3M)
- Hybrid approaches
  - Nested Grid Particle Scheme
  - Tree Codes
  - Tree Code Particle Mesh (TPM)
- Self Consistent Field (SCF), Smoothed-Particle Hydrodynamics (SPH), Symplectic etc.

Particle-Particle method

- Simplest. Adopts an all-pairs approach.
- State of the system at time `t` given by particle positions $x_i(t)$ and velocity $v_i(t)$ for `i=1` to `N`

$$\{x_i(t), v_i(t); i = 1, N\}$$

  – Steps:

1. Compute forces
2. Integrate equations of motion
3. Update time counter
   Each iteration updates $x_i(t)$ and $v_i(t)$ to compute $x_i(t + \Delta t)$ and $v_i(t + \Delta t)$

Particle-Particle Method

1. Compute forces

```
//initialize forces
for i=1 to N
  F_i = 0

//Accumulate forces
for i=1 to N-1
  for j=i+1 to N
      F_i = F_i + F_ij
      F_j = F_j - F_ij
```

F_{ij} is the force on particle i due to particle j

Typically: $F_i = F_{external} + F_{nearest\_neighbor} + F_{N\text{-}Body}$

Particle-Particle Method

2. Integrate equations of motion

```
for i=1 to N
```

$$v_i^{new} = v_i^{old} + \frac{F_i}{m_i}\,\Delta t$$ `//using a=F/m and v=u+at`

$$x_i^{new} = x_i^{old} + v_i\,\Delta t$$

3. Update time counter

$$t^{new} = t^{old} + \Delta t$$

Particle-Particle Method

```
t=0
while(t<t^final) {
//initialize forces
      for i=1 to N
        F_i = 0
//Accumulate forces
      for i=1 to N-1
        for j=i+1 to N
          F[i] = F[i] + F_ij
          F[j] = F[j] - F_ij
//Integrate equations of motion
      for i=1 to N
        v_i^new = v_i^old + (F_i / m_i) Δt  //using a=F/m and v=u+at
        x_i^new = x_i^old + v_i Δt
// Update time counter
      t = t + Δt
}
```

Particle-Particle Method

- Costs (CPU operations)?

```
t=0
while(t<t^final) {
//initialize forces
        for i=1 to N
          F_i = 0
//Accumulate forces
        for i=1 to N-1
          for j=i+1 to N
            F[i] = F[i] + F_ij
            F[j] = F[j] - F_ij
//Integrate equations of motion
        for i=1 to N
          v_i^new = v_i^old + (F_i/m_i) Δt  //using a=F/m and v=u+at
          x_i^new = x_i^old + v_i Δt
// Update time counter
        t = t + Δt
}
```

Particle-Particle Method

- Experimental results (then):
  - Intel Delta = 1992 supercomputer, 512 Intel i860s
  - 17 million particles, 600 time steps, 24 hours elapsed time

    M. Warren and J. Salmon

    Gordon Bell Prize at Supercomputing 1992
  - Sustained 5.2 Gigaflops = 44K Flops/particle/time step
  - 1% accuracy
  - *Direct method (17 Flops/particle/time step) at 5.2 Gflops would have taken 18 years, 6570 times longer*

Courtesy: CS267 Lecture 24 https://sites.google.com/lbl.gov/cs267-spr2019/

Particle-Particle Method

- Experimental results (now):

  Vortex particle simulation of turbulence
  - Cluster of 256 NVIDIA GeForce 8800 GPUs
  - 16.8 million particles
    - T. Hamada, R. Yokota, K. Nitadori. T. Narumi, K. Yasoki et al
    - *Gordon Bell Prize for Price/Performance at Supercomputing 2009*
  - Sustained 20 Teraflops, or $8/Gigaflop

Courtesy: CS267 Lecture 24 https://sites.google.com/lbl.gov/cs267-spr2019/

Particle-Particle (PP) Method

- Discussion
  - Simple/trivial to program
  - High computational cost
    - Useful when number of particles are small (few thousands) and
    - We are interested in close-range dynamics when the particles in the range contribute *significantly* to forces
    - Constant time step must be replaced with variable time steps and numerical integration schemes for close-range interactions

N-Body Simulation

- All-pairs Method
  - Naïve approach. Compute *all pair-wise interactions*
- Hierarchical Methods
  - Optimize. Reduce the number of pair-wise force calculations. How? dependence on ‘distant’ particle(s) can be *compressed*
  - Examples:
    - Barnes-Hut
    - Fast Multipole Method

Tree Codes

$$F_i = F_{external} + F_{nearest\_neighbor} + F_{N\text{-}Body}$$

- $F_{external}$ can be computed for each body independently. O(N)
- $F_{nearest\_neighbor}$ involve computations corresponding to few nearest neighbors. O(N)
- $F_{N\text{-}Body}$ require all-to-all computations. Most expensive. $O(N^2)$ if computed using all-pairs approach.

```
for(i = 1 to N)
```

$F_i = \sum_{i \neq j} F_{ij}$ F_{ij}= force on i from j

F_{ij} =c*v/||v||^3 in 3D, F_{ij} = c*v/||v||^2 in 2D

v = vector from particle i to particle j , ||v|| = length of v, c = product of masses or charges

We can do better.

Tree Codes: Divide-Conquer Approach

- Consider computing force on earth due to all celestial bodies
  - ➢ Look at the night sky. Number of terms in $\sum_{i \neq j} F_{ij}$ is greater than the number of visible stars
  - ➢ One "star" could really be the Andromeda galaxy, which contains billions of real stars. *Seems like a lot more work than we thought …*
  - – Idea: Ok to approximate all stars in Andromeda by a single point at its center of mass (CM) with same total mass (TM)

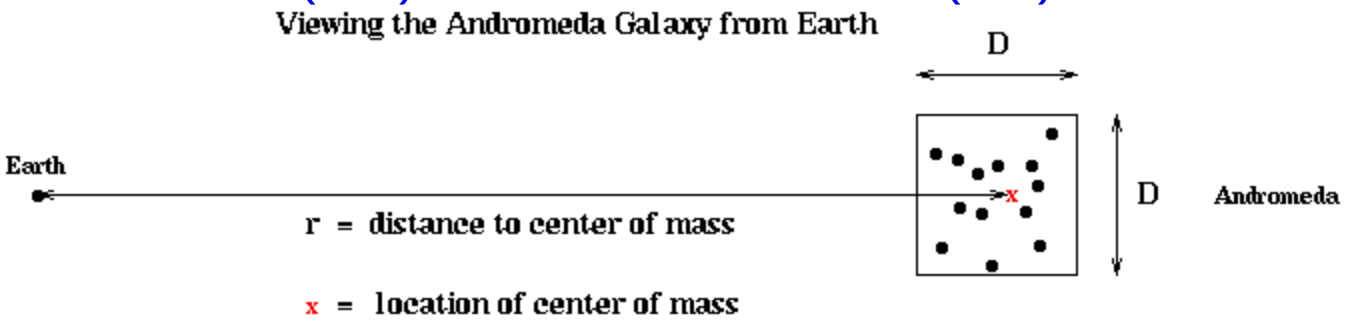

  - – Require that `D/r` be "small enough" (`D` = size of box containing Andromeda , r = distance of CM to Earth).

Idea is not new. Newton approximated earth and falling apple by CM

Slide contents based on: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Tree Codes: Divide-Conquer Approach

- New idea: recursively divide the box.
  - If you are in Andromeda, Milky Way (the galaxy we are part of) could appear like a white dot. So, can be approximated by a point mass.
  - Within Andromeda, picture repeats itself
    - As long as `D1/r1` is small enough, stars inside smaller box can be replaced by their CM to compute the force on Vulcan
    - If you are on Vulcan, another solar system in Andromeda can be a white dot.
    - Boxes nest in boxes recursively

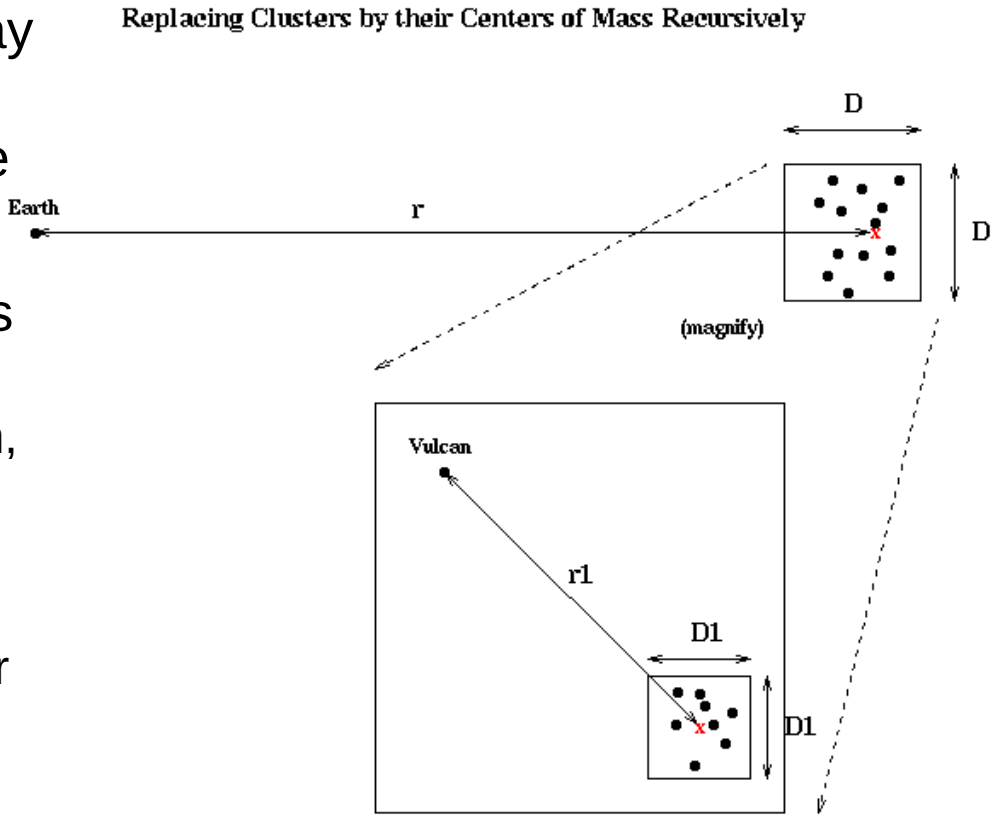

Slide contents based on: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Tree Codes: Divide-Conquer Approach

- Data structures needed:
  - Quad-trees
  - Octrees

Background – metric trees

e.g. K-dimensional (kd-), Vantage Point (vp-), quad-trees, octrees, ball-trees

2-dimensional space of points

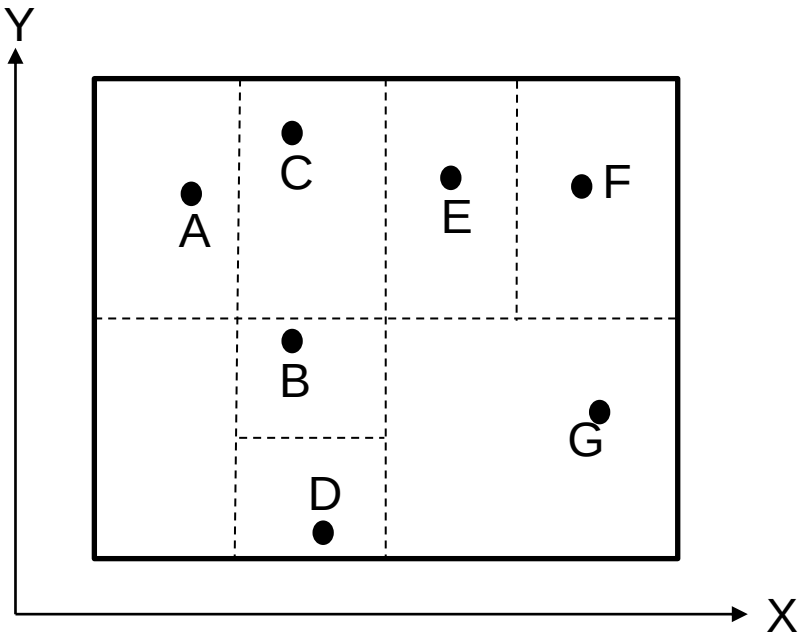

Binary kd-tree, 1 point /leaf cell

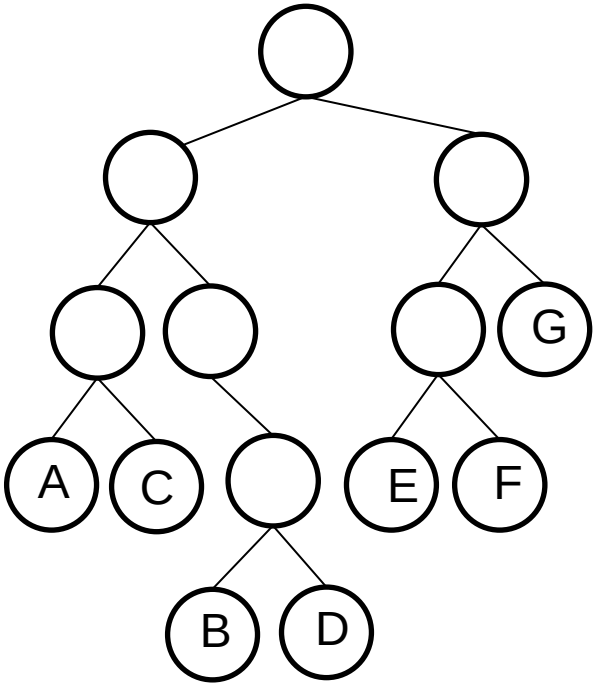

Background - metric trees

Typical use: traverse the tree (often repeatedly), truncate the traversal at some intermediate node if <u>a domain-specific criteria</u> is not met.

E.g. Does the distance from CM to me < D/r?

Input points = {1, 2, … , N} $\in \mathbb{R}^K$

Cost ???

Kd-tree

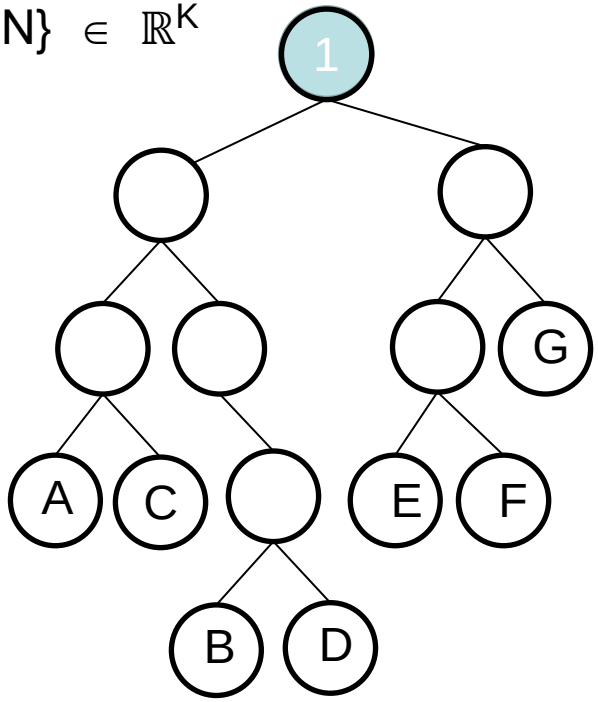

Quad Tree

- Data structure to subdivide the plane
  - Nodes can contain coordinates of center of box, side length.
  - Eventually also coordinates of CM, total mass, etc.
- In a complete quad tree, each non-leaf node has 4 children

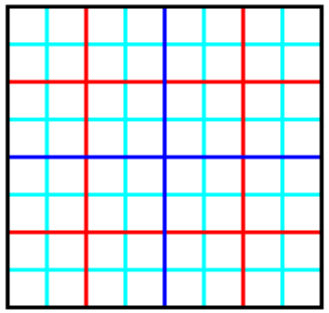

A Complete Quadtree with 4 Levels

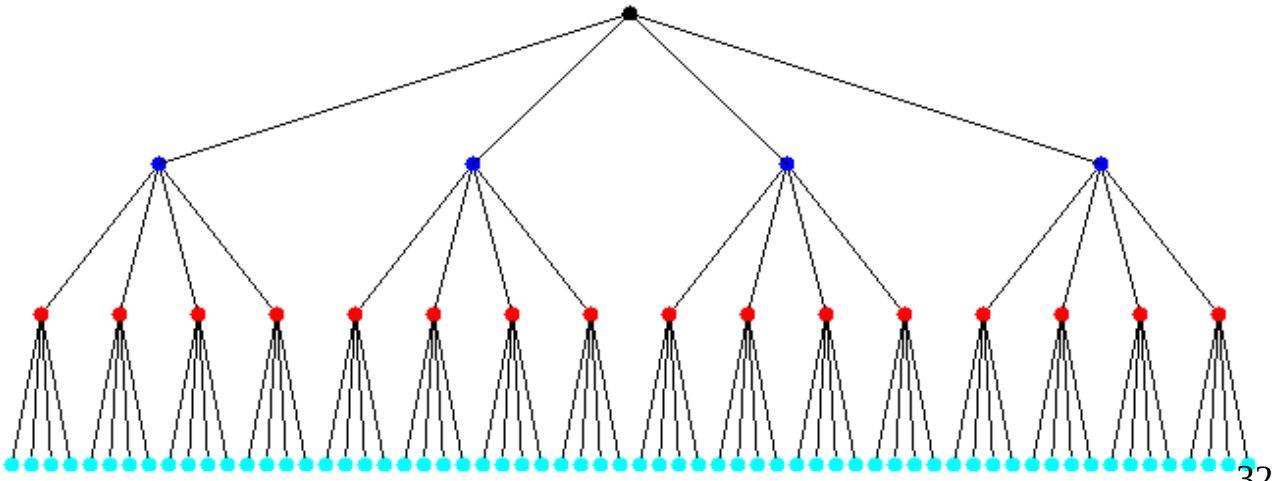

Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Octree or Oct Tree

- Similar data structure for subdividing 3D space

2 Levels of an Octree

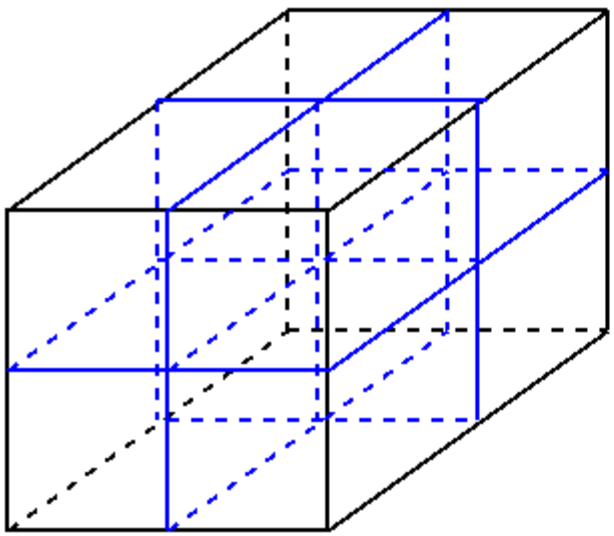

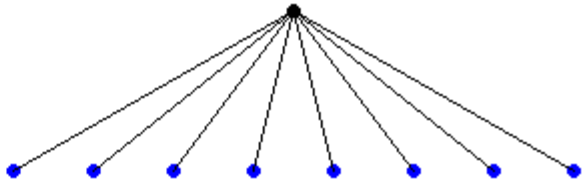

Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Using Quad Tree and Octree

- Begin by constructing a tree to hold all the particles
  - Interesting cases have nonuniformly distributed particles
  - In a complete tree most nodes would be empty, a waste of space and time
  - Adaptive Quad (Oct) Tree only subdivides space where particles are located
- For each particle, traverse the tree to compute force on it

Slide contents based on: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Using Quad Tree and Octree

Adaptive quadtree where no square contains more than 1 particle

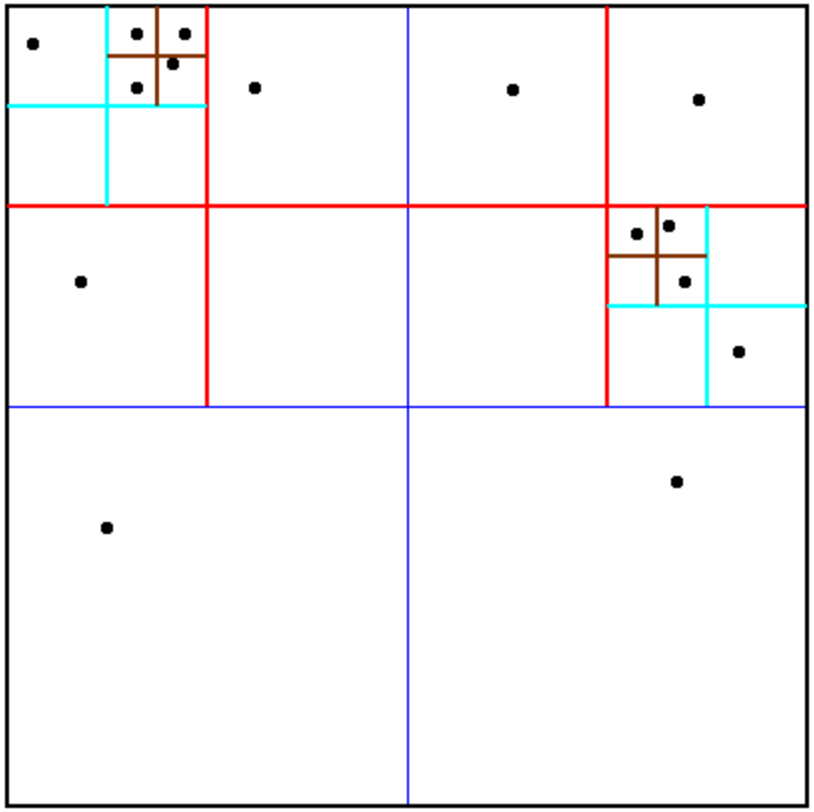

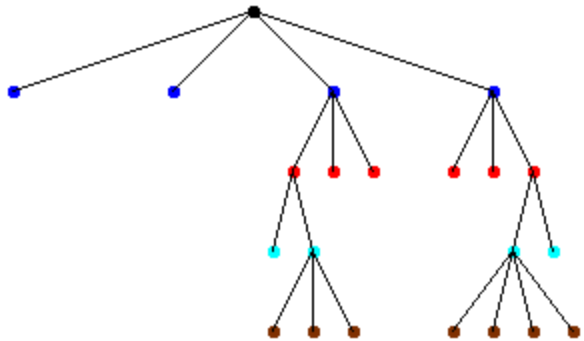

Child nodes enumerated counterclockwise from SW corner, empty ones excluded

- In practice, have q>1 particles/square; tuning parameter (code to build data structure on hidden slide)

Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Adaptive Quad Tree

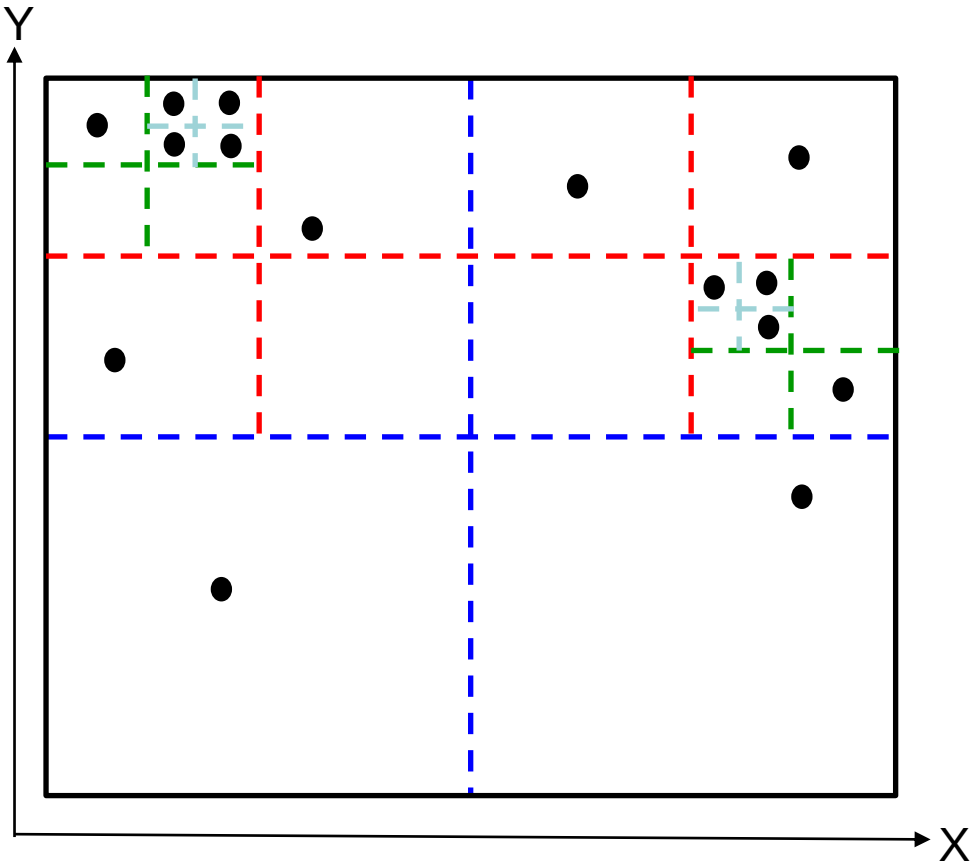

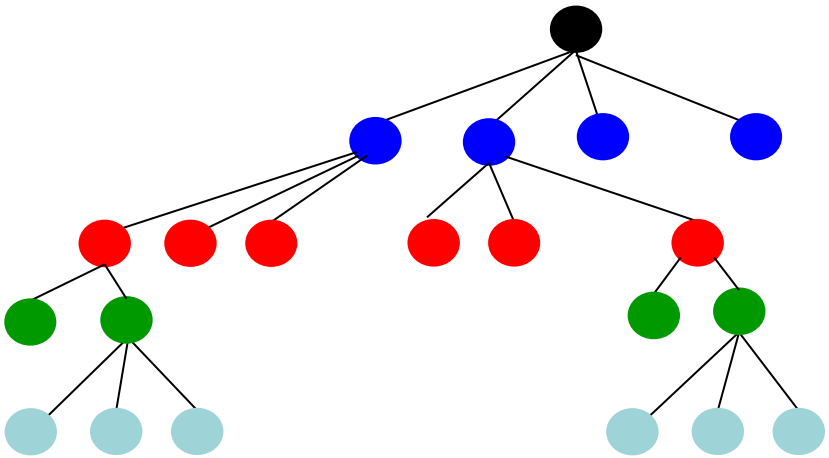

- In practice, #particles/square > 1. tuning parameter
- Child nodes numbered as per *Z-order numbering*

Adaptive Quad Tree Construction

```
Procedure Quad_Tree_Build
  Quad_Tree = {emtpy}
  for j = 1 to N                          … loop over all N particles
     Quad_Tree_Insert(j, root)            … insert particle j in QuadTree
  endfor
  …   At this point, each leaf of Quad_Tree will have 0 or 1 particles
  …   There will be 0 particles when some sibling has 1
  Traverse the Quad_Tree eliminating empty leaves  … via, say Breadth First Search

Procedure Quad_Tree_Insert(j, n)  … Try to insert particle j at node n in Quad_Tree
  if n an internal node             … n has 4 children
     - determine which child c of node n contains particle j
     - Quad_Tree_Insert(j, c)
  else if n contains 1 particle   …  n is a leaf
     - add n's 4 children to the Quad_Tree
     - move the particle already in n into the child containing it
     - let c be the child of n containing j
     - Quad_Tree_Insert(j, c)
  else                              …  n empty
     - store particle j in node n
  end
```


Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Adaptive Quad Tree Construction – Cost?

```
Procedure Quad_Tree_Build
   Quad_Tree = {emtpy}                          ≤ N *max cost of Quad_Tree_Insert
   for j = 1 to N                          … loop over all N particles
       Quad_Tree_Insert(j, root)           … insert particle j in QuadTree
   endfor
   …   At this point, each leaf of Quad_Tree will have 0 or 1 particles
   …   There will be 0 particles when some sibling has 1
   Traverse the Quad_Tree eliminating empty leaves  … via, say Breadth First Search
```

```
Procedure Quad_Tree_Insert(j, n) … Try to insert particle j at node n in Quad_Tree
   if n an internal node            … n has 4 children
      - determine which child c of node n contains particle j
      - Quad_Tree_Insert(j, c)
  else if n contains 1 particle   …  n is a leaf
      - add n's 4 children to the Quad_Tree
      - move the particle already in n into the child containing it
      - let c be the child of n containing j
      - Quad_Tree_Insert(j, c)                  ≤ max depth of Quad Tree
   else                             …  n empty
      - store particle j in node n
   end
```


Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Adaptive Quad Tree Construction – Cost?

- Max Depth of Tree:
  - For uniformly distributed points?
  - For arbitrarily distributed points?

- Total Cost = ?

Adaptive Quad Tree Construction – Cost?

- Max Depth of Tree:
  - For uniformly distributed points? = O(log N)
  - For arbitrarily distributed points? = O(bN)
    - b is number bits used to represent the coordinates

- Total Cost = $O(b\,N)$ or $O(N * \log N)$

Barnes-Hut

- Simplest hierarchical method for N-Body simulation
  - "A Hierarchical O(n log n) force calculation algorithm" by J. Barnes and P. Hut, Nature, v. 324, December 1986
- Widely used in astrophysics
- Accuracy $\geq$ 1% (good when low accuracy is desired/acceptable. Often the case in astrophysics simulations.)

Barnes-Hut: Algorithm

(2D for simplicity)

1) Build the QuadTree using QuadTreeBuild
   … already described, cost = O(N log N) or O(b N)

2) For each node/subsquare in the QuadTree, compute the
   Center of Mass (CM) and total mass (TM) of all the particles it contains.

3) For each particle, traverse the QuadTree to compute the force on it,

Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Barnes-Hut: Algorithm (step 2)

Goal: Compute the Center of Mass (CM) and Total Mass (TM) of all the particles in each node of the QuadTree. (TM, CM) = Compute_Mass(root)

```
(TM, CM)  = Compute_Mass( n )    //compute the CM and TM of node n
  if n contains 1 particle
       //TM and CM are identical to the particle's mass and location
       store (TM, CM) at n
       return (TM, CM)
  else
     for each child c(j) of n  //j = 1,2,3,4
           ( TM(j), CM(j) ) = Compute_Mass( c(j) )
     endfor
     TM = TM(1) + TM(2) + TM(3) + TM(4)
     //the total mass is the sum of the children's masses
     CM = ( TM(1)*CM(1) + TM(2)*CM(2) + TM(3)*CM(3) + TM(4)*CM(4) ) / TM
     //the CM is the mass-weighted sum of the children's centers of mass
     store ( TM, CM ) at n
     return ( TM, CM )
  end if
```


Slide based on : CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Barnes-Hut: Algorithm (step 2 cost)

(2D for simplicity)

1) Build the QuadTree using QuadTreeBuild
   … already described, cost = O(N log N) or O(b N)

2) For each node/subsquare in the QuadTree, compute the
   Center of Mass (CM) and total mass (TM) of all the particles it contains.
   … cost = O(number of nodes in the tree) = O(N log N) or O(b N)

3) For each particle, traverse the QuadTree to compute the force on it,

Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Barnes-Hut: Algorithm (step 3)

Goal: Compute the force on each particle by traversing the tree. For each particle, use as few nodes as possible to compute force, subject to accuracy constraint.

- For each node = square, can approximate force on particles outside the node due to particles inside node by using the node's CM and TM
- This will be accurate enough if the node if "far away enough" from the particle
- Need criterion to decide if a node is far enough from a particle
  - **D = side length of node**
  - **r = distance from particle to CM of node**
  - **θ = user supplied error tolerance < 1**
  - **Use CM and TM to approximate force of node on box if D/r < θ**

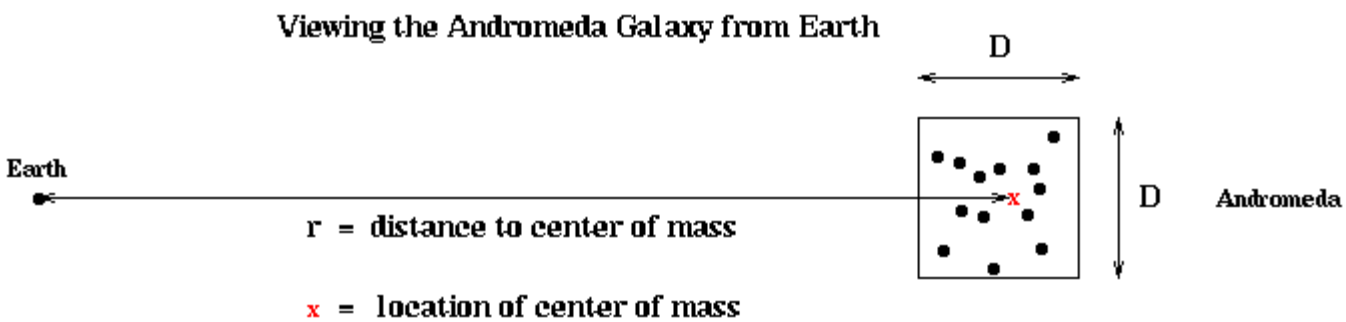

Slide based on : CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Barnes-Hut: Algorithm (step 3)

```
//for each particle, traverse the QuadTree to compute the force on it
for k = 1 to N
    f(k) = TreeForce( k, root )
    //compute force on particle k due to all particles inside root (except k)
endfor
function f = TreeForce( k, n )
    //compute force on particle k due to all particles inside node n (except k)
    f = 0
    if n contains one particle (not k)  //evaluate directly
        return f = force computed using direct formula
    else
        r = distance from particle k to CM of particles in n
        D = size of n
        if  D/r  < q    //ok to approximate by CM and TM
             return f = computed approximately using CM and TM
        else              //need to look inside node
             for each child c(j) of n //j=1,2,3,4
                  f = f + TreeForce ( k, c(j) )
             end for
             return f
        end if
    end if
```

Slide based on : CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Barnes-Hut: step 3 example

- Example: Assume $\theta \geq 1$. In practice $\theta < 1$.

What is the force on Point 1 due to all other points in the box with black-boundary?

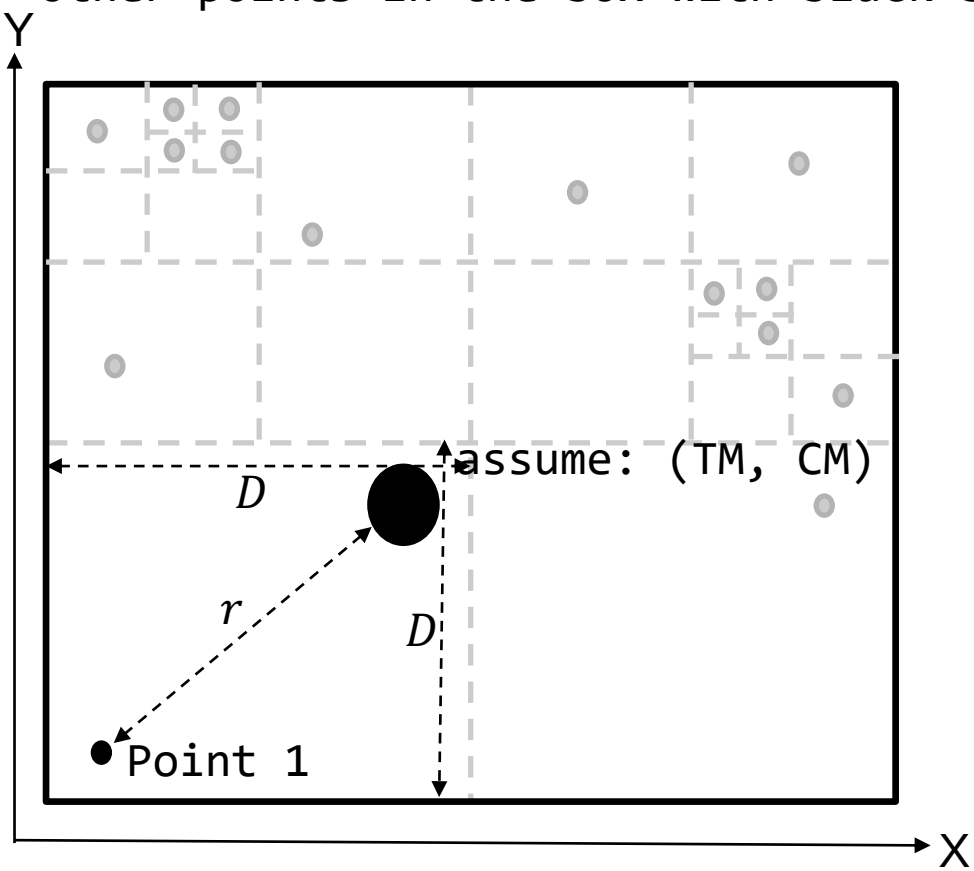

Point 1: is $D/r < \theta$?

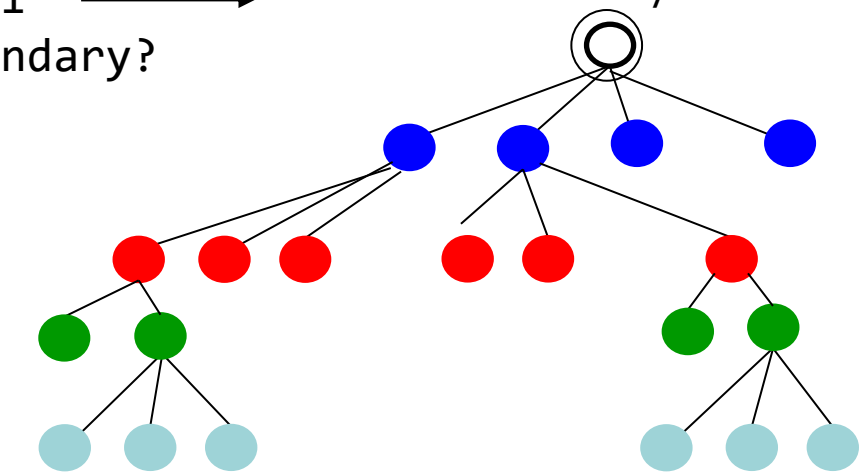

No. Compute force due to each child of the root node (i.e. particles in each quadrant of the square). Start with child 1: c(1).

Barnes-Hut: step 3 example

- Example: Assume $\theta \geq 1$. In practice $\theta < 1$.

What is the force on Point 1 due to all other points in the box with black-boundary?

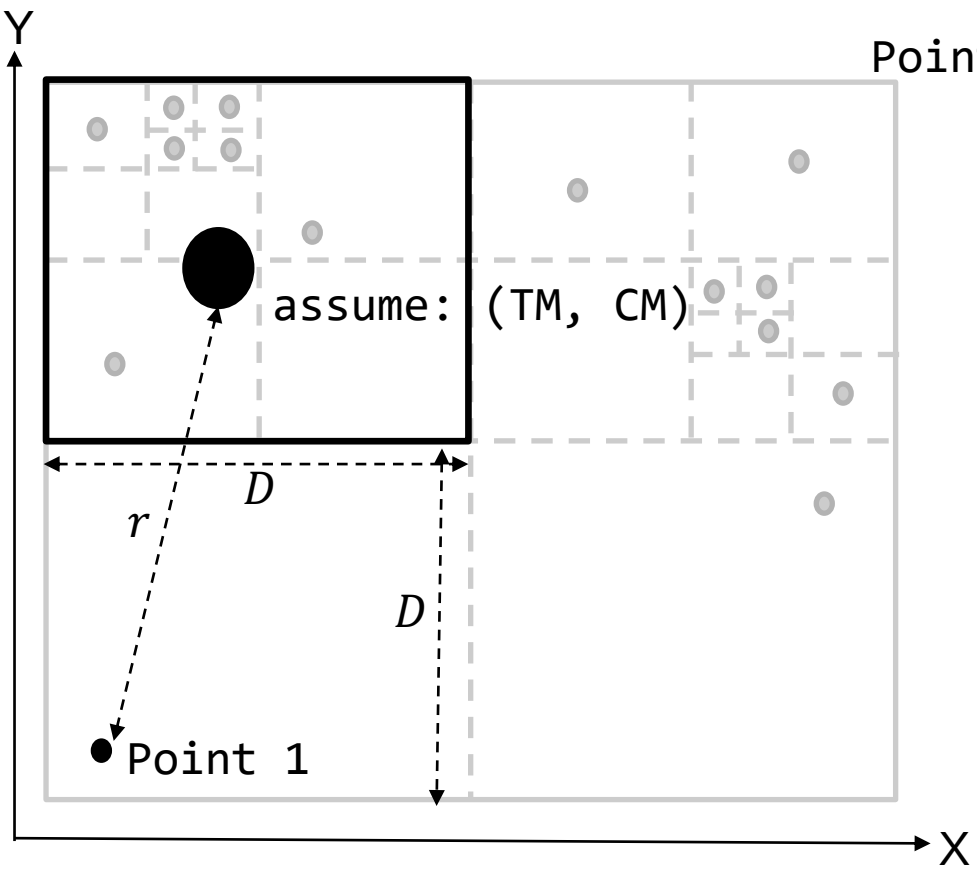

Point 1: is $D/r < \theta$?

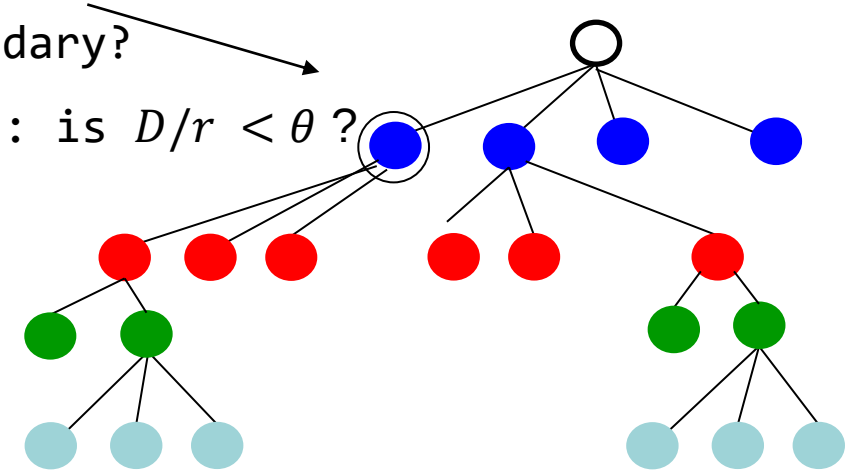

Yes. Approximate force due to each particle contained in the black-boundary box by the TM and CM of the box.

Barnes-Hut: step 3 example

- Example: Assume $\theta \geq 1$. In practice $\theta < 1$.

What is the force on Point 1 due to all other points in the box with black-boundary?

Point 1: is $D/r < \theta$?

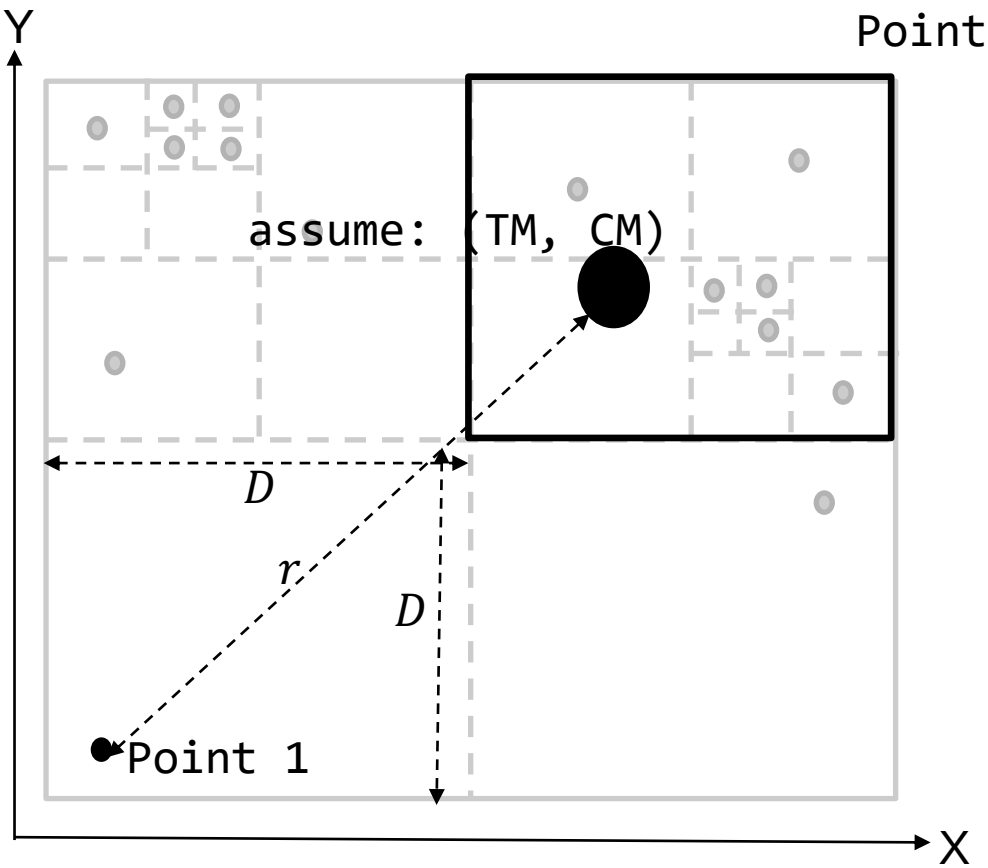

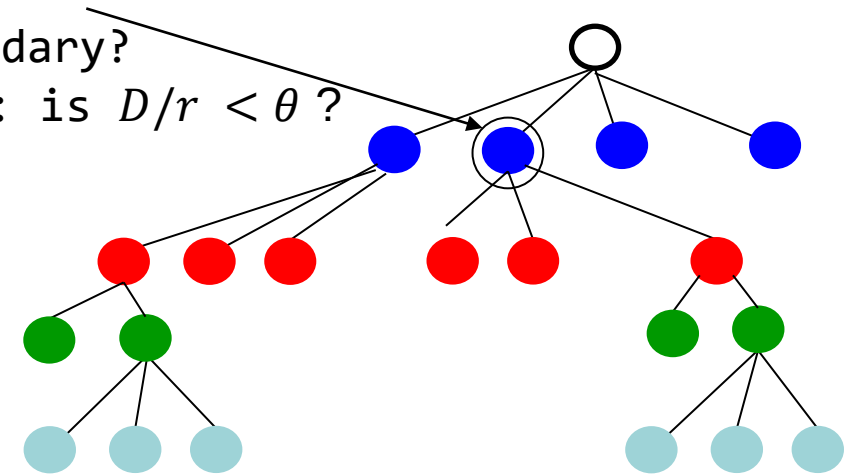

Yes. Approximate force due to each particle contained in the black-boundary box by the TM and CM of the box.

Barnes-Hut: step 3 example

- Example: Assume $\theta \geq 1$. In practice $\theta < 1$.

What is the force on Point 1 due to all other points in the box with black-boundary?

Point 1: is $D/r < \theta$?

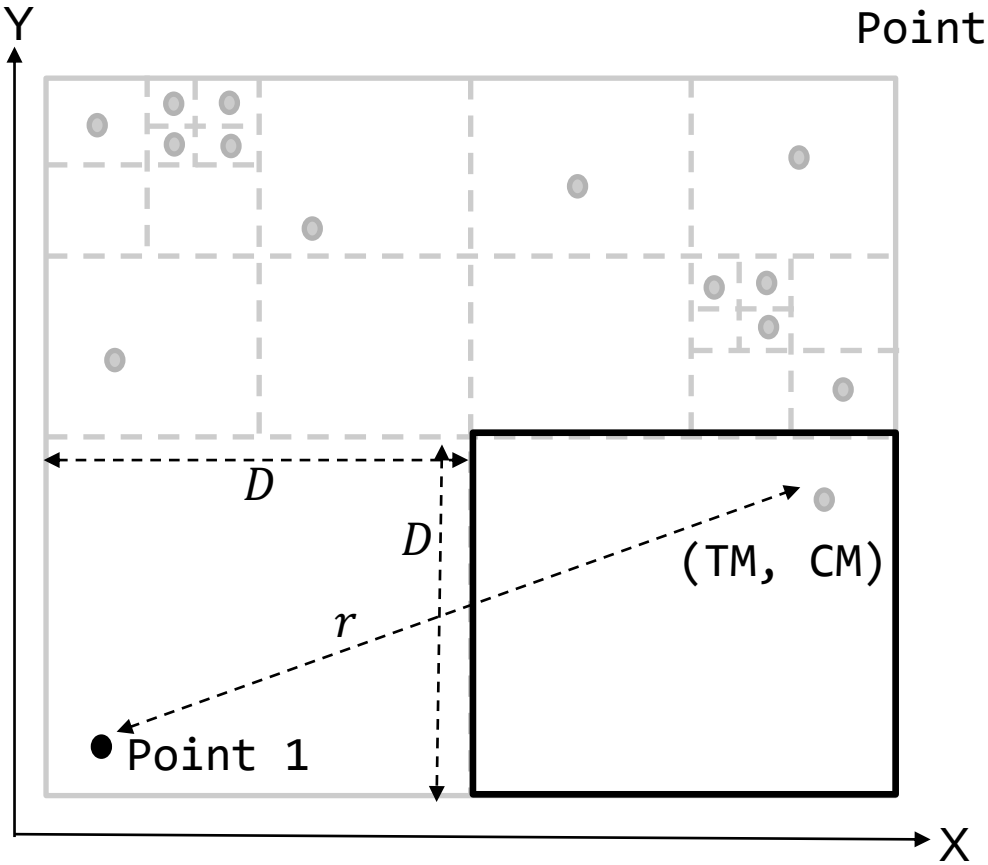

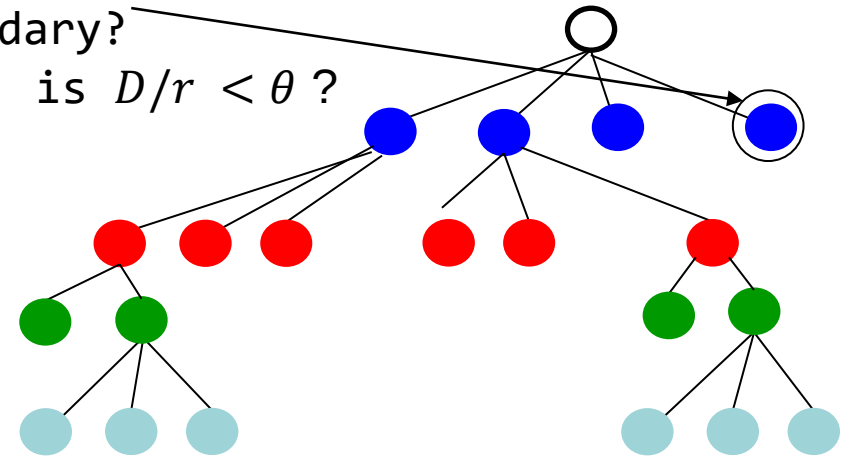

Contains 1 particle / leaf node. Compute force using direct formula.

Barnes-Hut: step 3 example

- Example: Assume $\theta \geq 1$. In practice $\theta < 1$.

What is the force on **Point 2** due to all other points in the box with black-boundary?

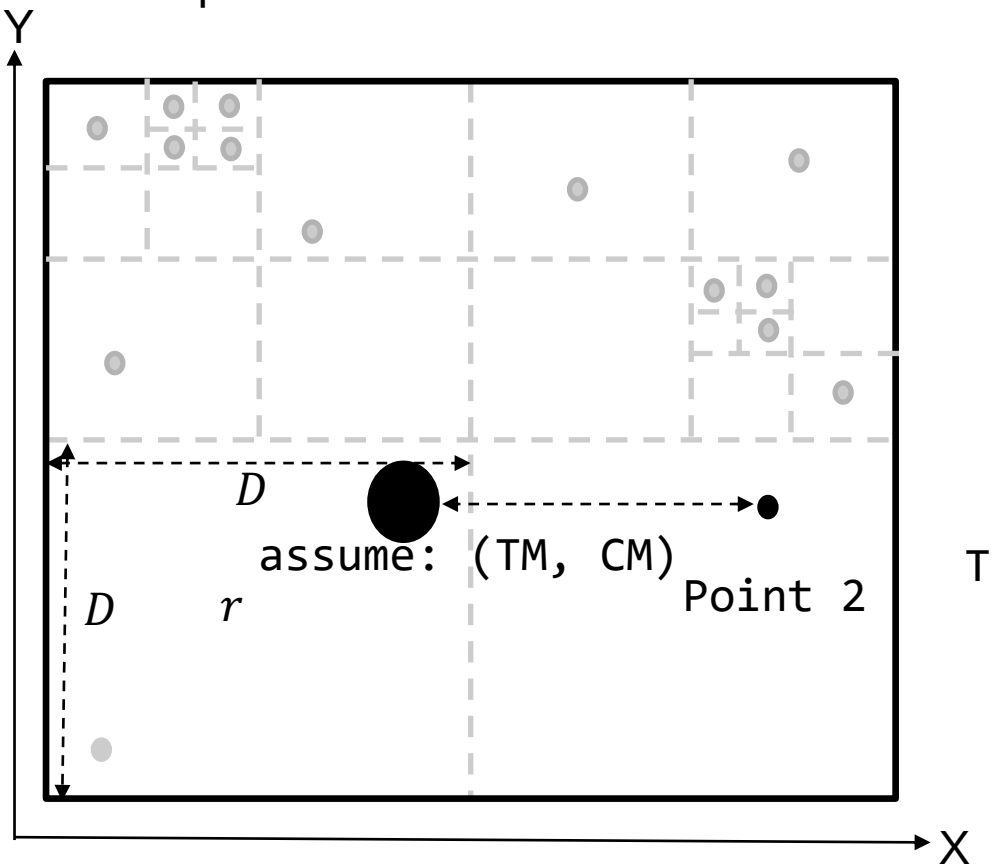

Point 2: is $D/r < \theta$?

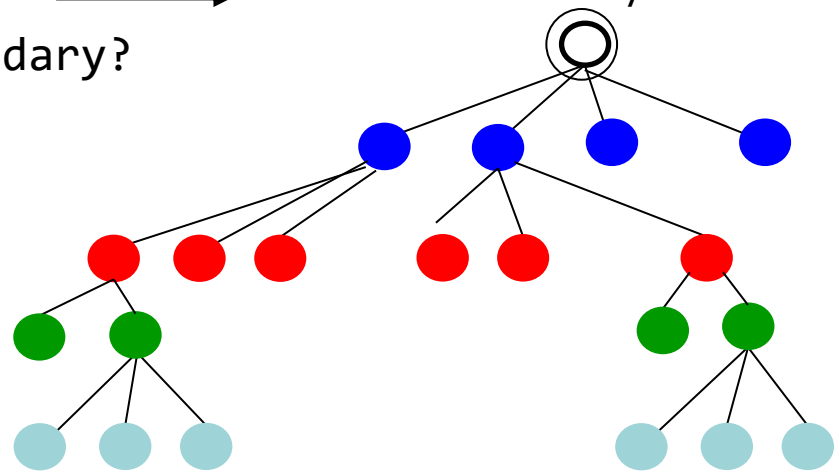

Traverse the tree for particle 2.

Barnes-Hut: Algorithm (step 3 cost)

- **Correctness** follows from recursive accumulation of force from each subtree
  - Each particle is accounted for exactly once, whether it is in a leaf or other node
- **Complexity** analysis
  - **Cost of TreeForce(k, root) = O(depth of leaf containing k in the QuadTree)**
  - Proof by Example (for θ>1):
    - For each undivided node = square, (except one containing k), D/r < 1 < θ
    - There are at most 3 undivided nodes at each level of the QuadTree.
      - There is O(1) work per node
      - Cost = O(level of k)

Total cost = O(Σ_k level of k) = O(N log N)

Strongly depends on θ

Sample Barnes–Hut Force calculation
For particle in lower right corner
Assuming theta > 1

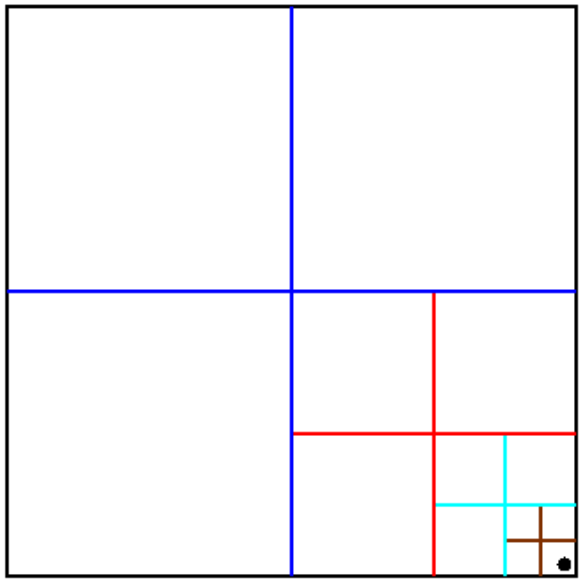

Slide based on : CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

Barnes-Hut: Algorithm (step 3 cost)

(2D for simplicity)

1) Build the QuadTree using QuadTreeBuild
   … already described, cost = O(N log N) or O(b N)

2) For each node/subsquare in the QuadTree, compute the Center of Mass (CM) and total mass (TM) of all the particles it contains.
   … cost = O(number of nodes in the tree) = O(N log N) or O(b N)

3) For each particle, traverse the QuadTree to compute the force on it,
   … cost depends on accuracy desired (θ) but still O(N log N) or O(bN)

Slide courtesy: CS267 Lecture 24, https://sites.google.com/lbl.gov/cs267-spr2019/

N-Body Simulation: Big Picture

- Recall:

```
t=0
while(t<t^final) {
//initialize forces

//Accumulate forces
      BH(steps 1 to 3)

//Integrate equations of motion

//Update time counter
      t = t + Δt
}
```

Fast Multipole Method (FMM)

- Can we make the complexity independent of the accuracy parameter (θ) ? FMM achieves this.
  - "Rapid Solution of Integral Equations of Classical Potential Theory", V. Rokhlin, J. Comp. Phys. v. 60, 1985 and
  - **"A Fast Algorithm for Particle Simulations",** L. Greengard and V. Rokhlin, J. Comp. Phys. v. 73, 1987.
- Similar to BH:
  - uses QuadTree and the divide-conquer paradigm
- Different from BH:
  - Uses more than TM and CM information in a box. So, computation is expensive and accurate than BH.
  - The number of boxes evaluated is fixed for a given accuracy parameter
  - Computes potential and not the Force as in BH

Background: Potential

- Force on a particle at (x, y, z) due to a particle at origin

$\propto -\frac{(x,y,z)}{r^3}$ (This is called inverse-square law. Gravitational and electrostatic forces obey this.) where, $r = \sqrt{x^2 + y^2 + z^2}$

- Force is a vector. Potential is a scalar. Hence, potential is simple to deal with.

Potential $\Phi(x, y, z) = -\frac{1}{r}$

- Negative of the gradient of potential = force

$$-\nabla\Phi(x, y, z) = -\left(\frac{d}{dx}\left(-\frac{1}{r}\right), \frac{d}{dy}\left(-\frac{1}{r}\right), \frac{d}{dz}\left(-\frac{1}{r}\right)\right)$$

Background: Potential

- In 2D, potential $\Phi(x, y) = \log r$
- Suppose we have N points (at $z_1, z_2, \ldots, z_N$, where $z_i = (x_i, y_i)$) in a plane with masses $m_1, m_2, \ldots, m_N$ resp.

then, their potential at $z = (x, y)$ is given by:

$$\Phi(x, y) = \sum_{i=1}^{N} m_i \log\left(\sqrt{(x - x_i)^2 + (y - y_i)^2}\right)$$

Goal: *evaluate* $\Phi(x, y)$ *and its derivatives at N points* $(z_1, z_2, \ldots, z_N)$ *in O(N) time.*

FMM Algorithm

1. Build the quadtree containing all the points.
2. Traverse the quadtree from bottom to top, computing Outer(n) for each square n in the tree.
3. Traverse the quadtree from top to bottom, computing Inner(n) for each square in the tree.
4. For each leaf, add the contributions of nearest neighbors and particles in the leaf to Inner(n)

what is Outer(n) and Inner(n) ?

Well Separated Regions

- Compute the influence of all particles in source region (B) on every particle in target region (A)

  (assumption: A and B are well-separated)

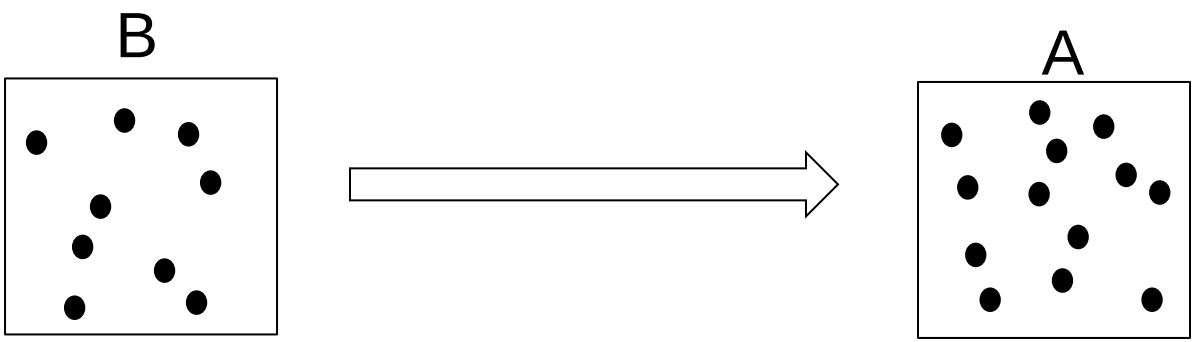

- At each point p_i in A, compute potential:

$$\Phi(x_i, y_i) = \sum_{p_j \in B} m_i \log |p_i - p_j|$$

$$i = 1 \; to \; N_A, \qquad j = 1 \; to \; N_B$$

- Cost: $O(N_A N_B)$

Well Separated Regions

- Approximate the potential at every particle in target region (A) by the potential at C_A

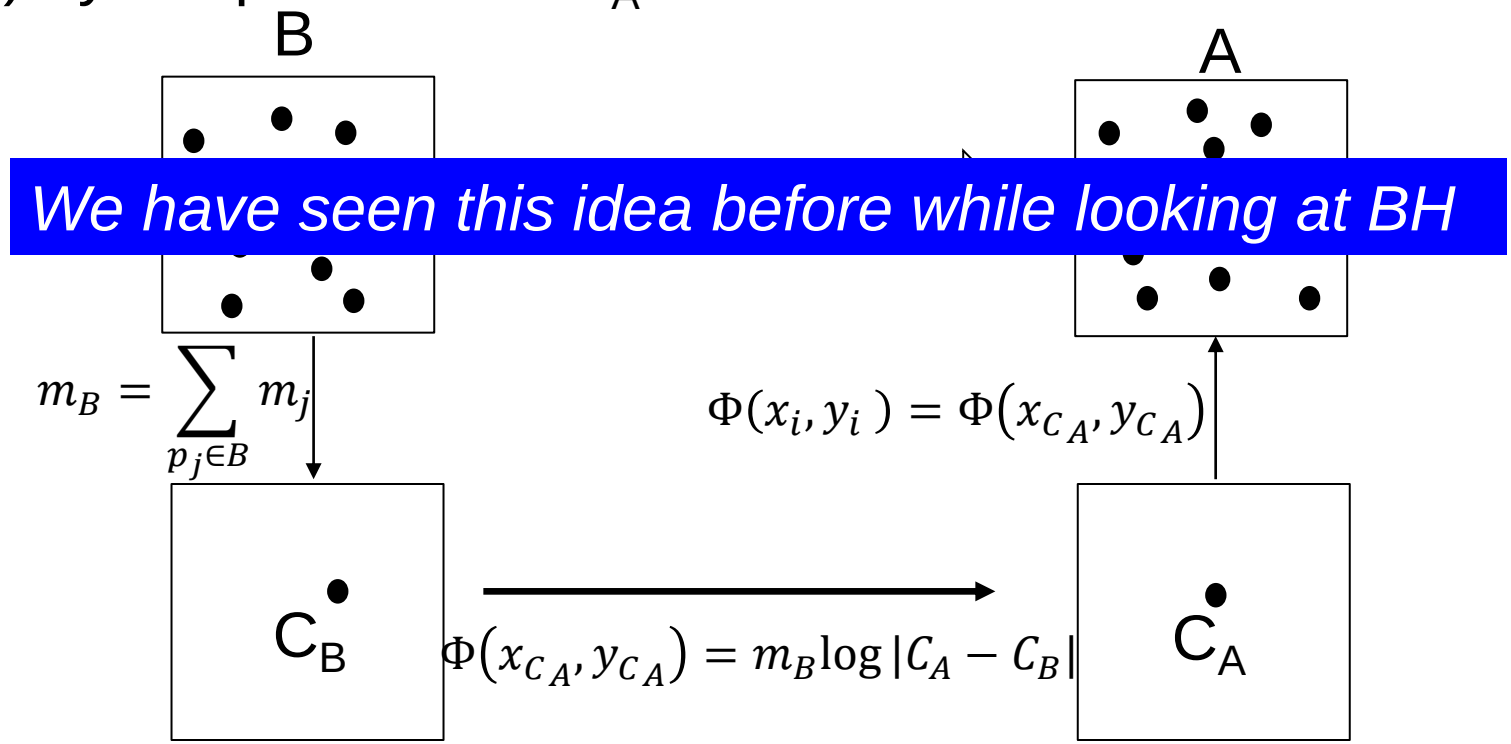

- Cost: $O(N_A + N_B)$

Hierarchical Decomposition

- In N-body simulation, every point serves as source as well as target. *How to identify source, target, well-separated regions?*
  - Partition the space recursively till every leaf box contains O(1) number of points

FMM Algorithm

1. Build the quadtree containing all the points.
2. Traverse the quadtree from bottom to top, computing Outer(n) for each square n in the tree.
3. Traverse the quadtree from top to bottom, computing Inner(n) for each square in the tree.
4. For each leaf, add the contributions of nearest neighbors and particles in the leaf to Inner(n)

what is Outer(n) and Inner(n) ?

Well Separated Regions

- Compute the influence of all particles in source region (B) on every particle in target region (A)

  (assumption: A and B are well-separated)

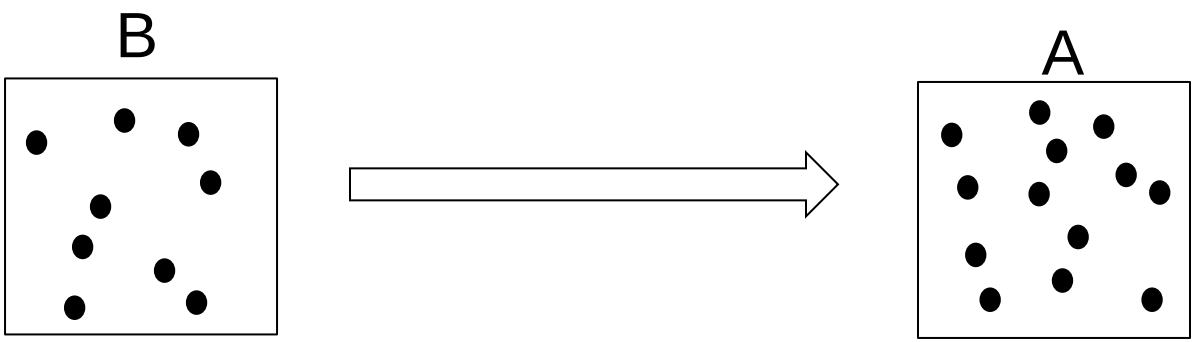

- At each point p_i in A, compute potential:

$$\Phi(x_i, y_i) = \sum_{p_j \in B} m_i \log |p_i - p_j|$$

$$i = 1 \; to \; N_A, \qquad j = 1 \; to \; N_B$$

- Cost: $O(N_A N_B)$

Well Separated Regions

- Compute the influence of all particles in source region (B) on every particle in target region (A)

$$\Phi\left(x_{p_i}, y_{p_i}\right) = \sum_{p_j \in B} m_i \log|p_i - p_j|\ , p_i \in A$$

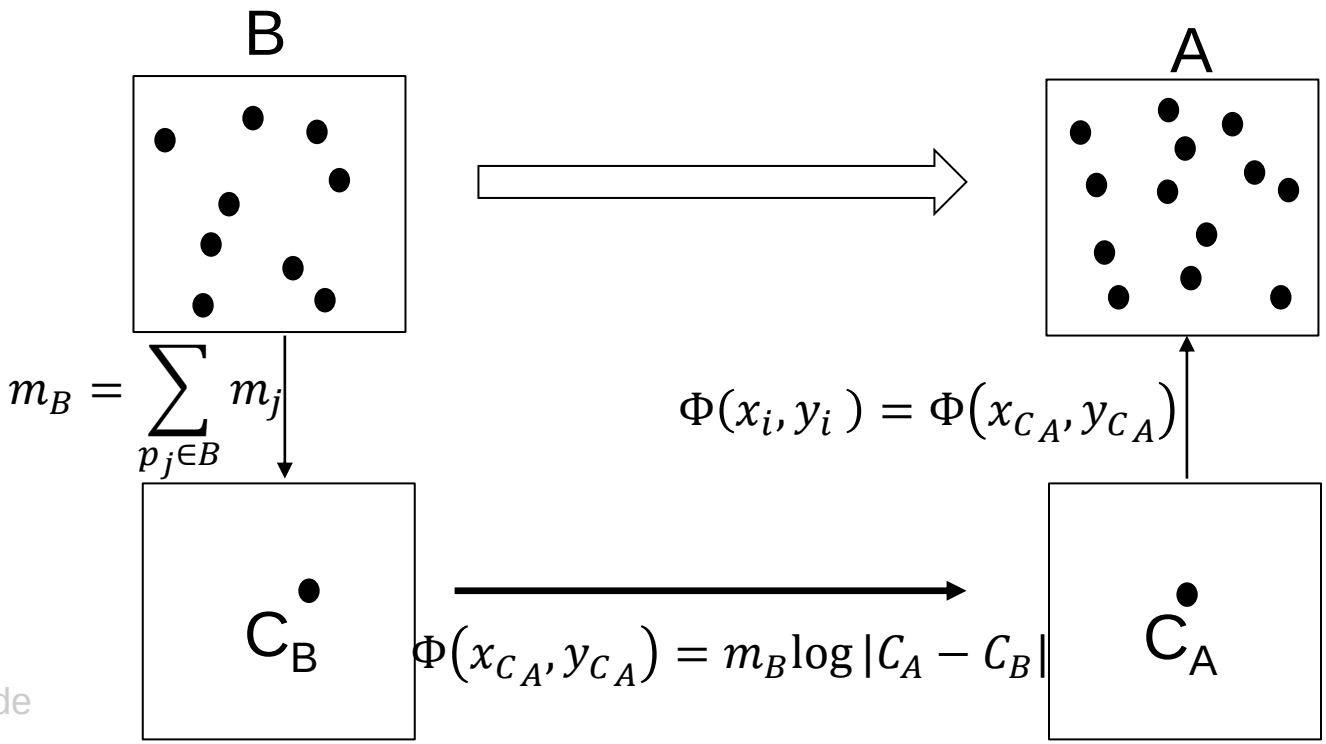

Applying the 3-step Approximation

- In N-body simulation every point serves as source as well as target.

How to identify source and target (boxes A and B in previous slide) i.e. well-separated regions?

Hierarchical decomposition

Hierarchical Decomposition

- *Level-0 decomposition*

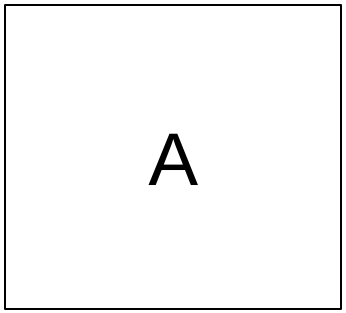

No well-separated boxes

- *Level-1 decomposition*

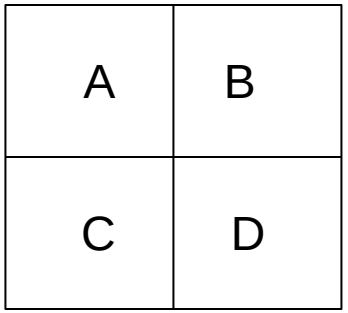

Hierarchical Decomposition

- *Level-2 decomposition*

N1	N2	N3	A1
N4	**B**	N5	A2
N6	N7	N8	A3
A7	A6	A5	A4

Well-separated from B

Can approximate the influence of points in `B` on points in `Ai` s

What do we do about **B**'s influence on Ni s?

Hierarchical Decomposition

- *Level-3 decomposition*

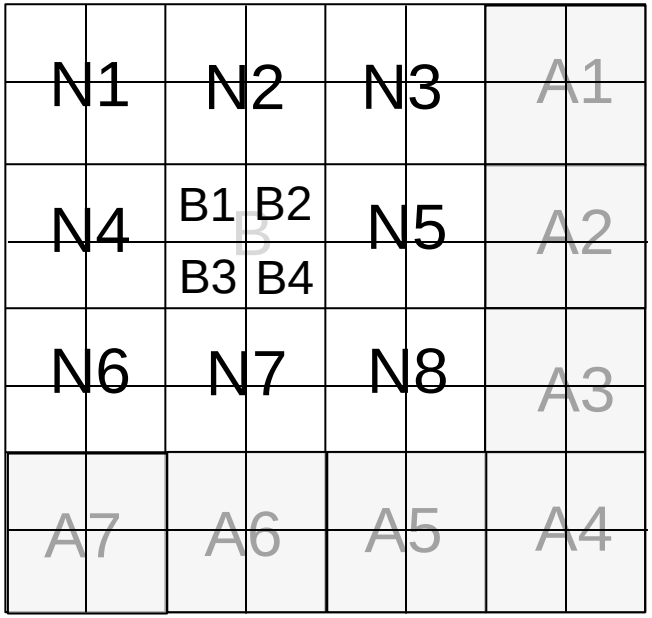

Influence of points in `Bi` s on those in `Ai` s already computed at the previous level (level-2)

Hierarchical Decomposition

- *Level-3 decomposition*

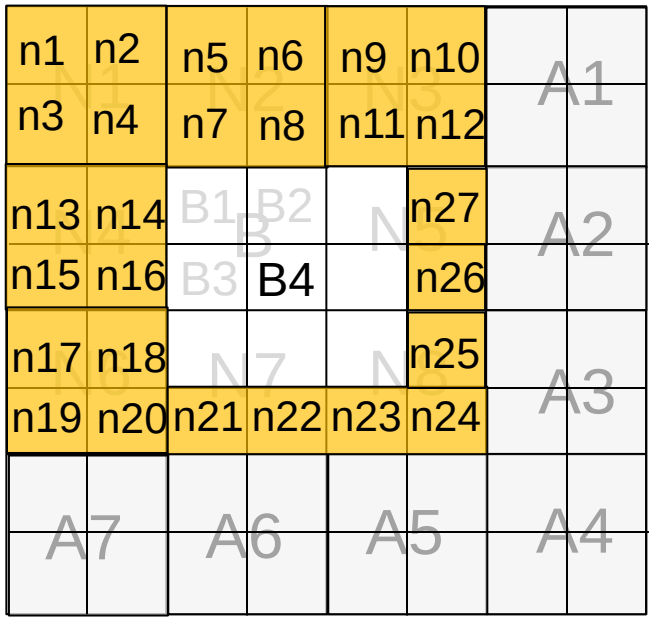

Influence of points in `Bi` s on those in `Ai` s already computed at the previous level (level-2)

Well-separated from B4

Influence of `B4's` points on `nx's` points can be approximated

`nx's` constitute the <u>interaction list</u> for `B4`.

What is the max size of interaction list? i.e. max number of `nx` *s that we can have for any* `Bi`*?*

Hierarchical Decomposition

- *Level-3 decomposition*

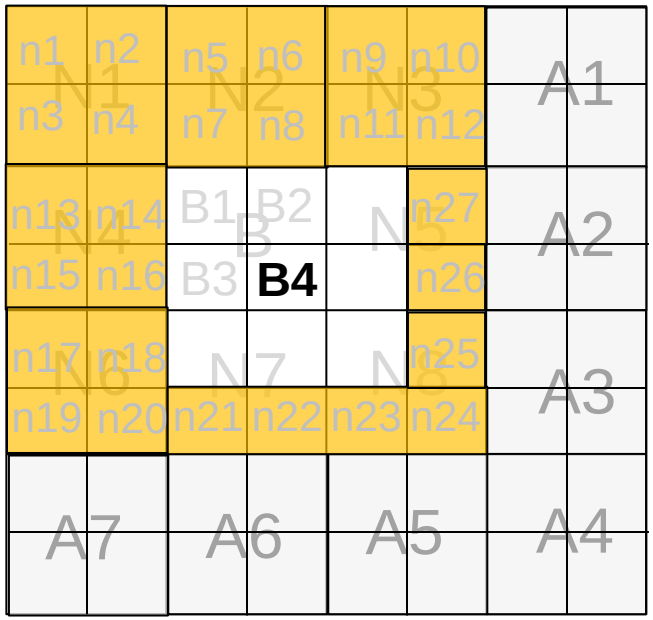

Influence of points in `Bi` s on those in `Ai` s already computed at the previous level (level-2)

Well-separated from B4

Influence of `B4's` points on `nx's` points can be approximated

What do we do about **B4**'s influence on its neighbors (white/unshaded boxes)?

Hierarchical Decomposition

- *Level-4 decomposition*

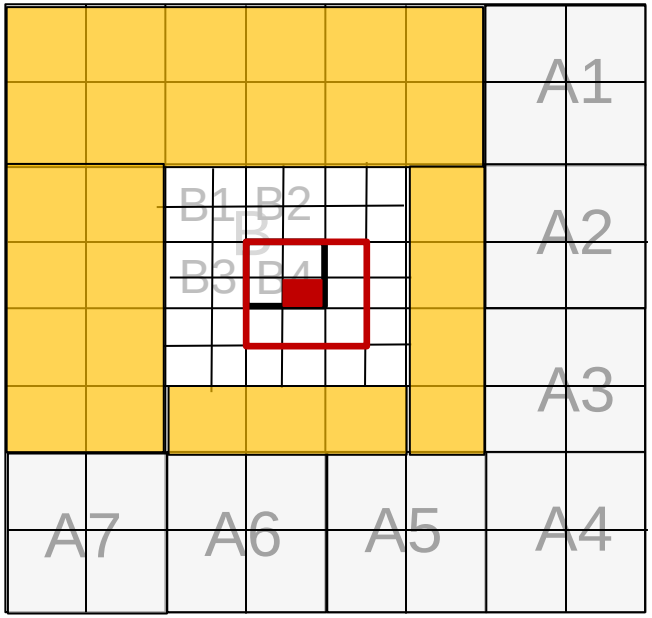

Any unshaded box outside ☐ can be the *target* for computing the influence of points in ■ (*source*)

1. Computing Potential for Well-Separated Regions

```
for level L=2 to last_level
  for each Box B at level L
        iList = GetInteractionList(B)
        for each well-separated box A in iList
              //Compute potential
              potential = m_B log|C_A − C_B|
              //Accumulate potential
              Φ(x_{C_A}, y_{C_A}) +=potential
```

Cost?

1. Computing Potential for Well-Separated Regions

```
for level L=2 to last_level
  for each Box B at level L
        iList = GetInteractionList(B)
        for each well-separated box A in iList
              //Compute potential
```

5. potential $= m_B \log|C_A - C_B|$

```
              //Accumulate potential
```

6. $\Phi(x_{C_A}, y_{C_A})$ +=potential

Prereqs: we need $\mathrm{m_B}, \mathrm{C_A}, \mathrm{C_B}$ details. (step 0)

2. Assigning Potential to Points

```
for each Box A at level L=0 to last_level
```

$\Phi_{p_i} = \Phi_{p_i} + \Phi_{C_A}$ (where $p_i \in A$ and C_A is A's CM)

Cost?

3. Assigning Potential to Points (last level)

```
1. for each Box B at last_level
```

2. $\Phi_{p_i} = \Phi_{p_i} + \sum_{p_j \in Neighbors(B)} m_B \log |p_i - p_j|$ (where $p_i \in B$)

Cost?

0. Computing Prereqs

```
for each Box B at level L=0 to last_level
    m_B = Σ_{p_j ∈ B} m_j
    //similarly compute C_B
```

Cost?

Total Cost (steps 0 + 1 + 2 + 3)

```
O(N log N) + O (N) + O(N log N) + O(N)
```

Can we do better?

0'. Computing Prereqs

- Traverse the tree bottom up instead of top-down

```
for each Box B starting from last_level to L=0
    if B is a leaf box
        m_B = Σ_{p_j ∈ B} m_j
    else
        m_B = m_{B_1} + m_{B_2} + m_{B_3} + m_{B_4}
        //B_1-B_4 are children of B
```

Cost?

2'. Assigning Potential to Points

```
for each Box A at level L=0 to last_level
    if A is a leaf box
```

$\Phi_{p_i} = \Phi_{p_i} + \Phi_{C_A}$ (where $p_i \in A$ and C_A is A's CM)

```
    else
```

$\Phi_{A_1} = \Phi_{A_1} + \Phi_A$

$\Phi_{A_2} = \Phi_{A_2} + \Phi_A$

$\Phi_{A_3} = \Phi_{A_3} + \Phi_A$

$\Phi_{A_4} = \Phi_{A_4} + \Phi_A$

//A_1-A_4 are children of A

Cost?

Total Cost (steps 0’ + 1 + 2’ + 3)

```
O(N) + O (N) + O(N) + O(N)
```

Problem: low accuracy if source (A) and target (B) are not far away from each other

Solution: more accurate representations for m_B and $\Phi\left(x_{C_A}, y_{C_A}\right)$

Multipole expansion

- Like a Taylor series expansion that is accurate when $x^2 + y^2$ is large (x, y are cartesian coordinates of the point)
- For a quadtree box B centered at $\left(x_{C_B}, y_{C_B}\right)$, we compute and store the terms: $\{\boldsymbol{m_B}, \boldsymbol{\alpha_1}, \boldsymbol{\alpha_2}, \ldots, \boldsymbol{\alpha_p}, \boldsymbol{z_{C_B}}\}$

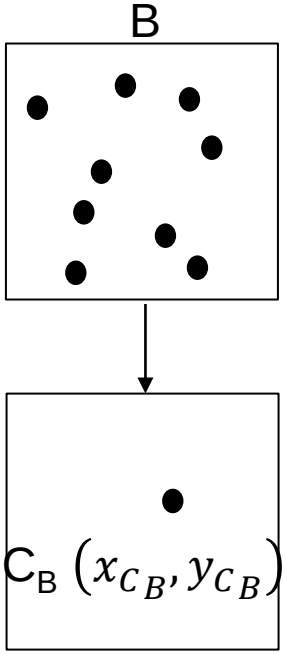

$$\alpha_j = \sum_{i=1}^{N_B} m_i \left(\frac{z_i^j}{j}\right)$$

$z_i$means$|z_i| = |(x_i, y_i)|$

Multipole expansion

- We approximate the potential at point z due to B by:

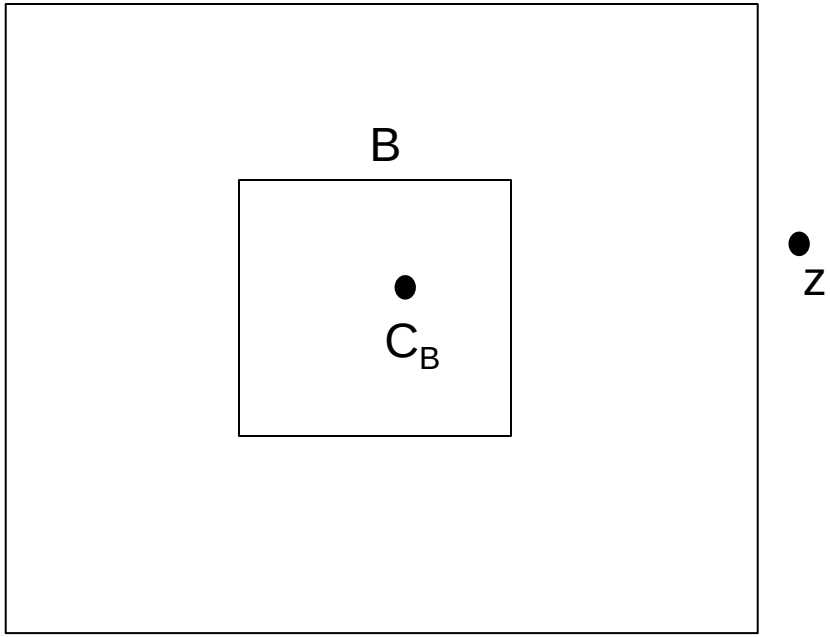

$$\Phi(x_z, y_z) = m_B \log(z - C_B) + \frac{\alpha_1}{z - C_B} + \frac{\alpha_2}{(z - C_B)^2} + \cdots + \frac{\alpha_p}{(z - C_B)^p}$$

- Because $\{\boldsymbol{m_B}, \boldsymbol{\alpha_1}, \boldsymbol{\alpha_2}, \ldots, \boldsymbol{\alpha_p}, \boldsymbol{z_{C_B}}\}$ is used to compute potential outside B, it is called <u>outer expansion</u>

Multipole expansion

- Similarly, we have the <u>inner expansion</u> $\{m_B, \beta_1, \beta_2, \ldots, \beta_p, z_{C_B}\}$

for computing the potential inside the Box due to all other points outside the box

- Computing outer expansions starts from leaf nodes and proceeds upwards in the tree.
- Computing inner expansions starts from root node and proceeds downwards in the tree.

3-Step Approximation (accurate)

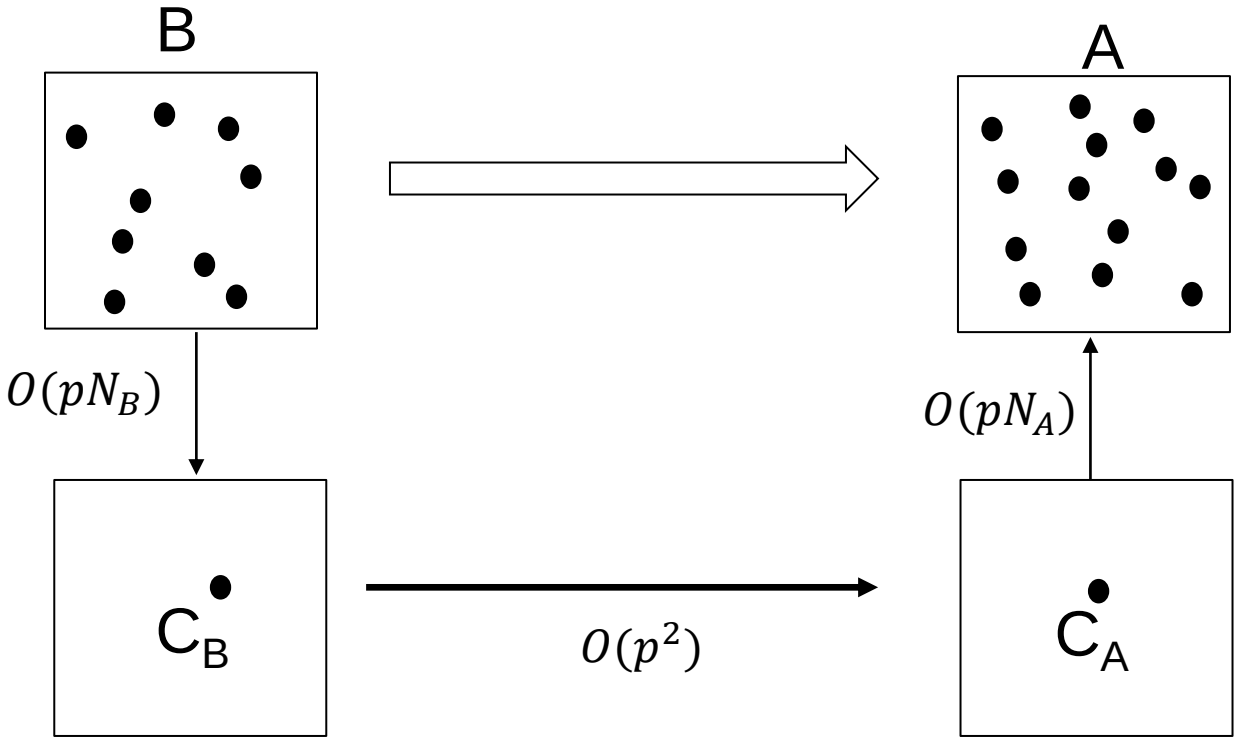

FMM Algorithm

1. Build the quadtree containing all the points.
2. Traverse the quadtree from bottom to top, computing Outer(n) for each square n in the tree.
3. Traverse the quadtree from top to bottom, computing Inner(n) for each square in the tree.
4. For each leaf, add the contributions of nearest neighbors and particles in the leaf to Inner(n)

Multipole expansion

- How to obtain the expression for alpha, beta ?
- What is the value of p?
- How to compute alpha and beta?

- Further reading:
  https://people.eecs.berkeley.edu/~demmel/cs267/lecture27/lecture27.html