Today

- Barnes-Hut

  $N^2 \to$ lower

  compromise cplx.
- FMM
- Direct solve

Announcements

- HW3
- Grading

Barnes-Hut: Putting Multipole Expansions to Work

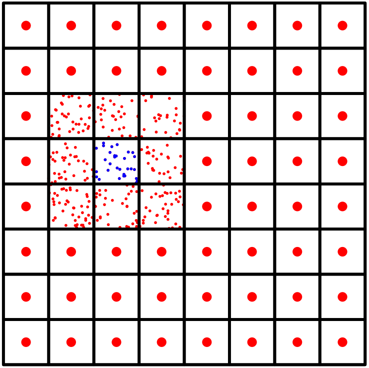

(Figure credit: G. Martinsson)

Barnes-Hut: Accuracy

With this computational outline, what's the accuracy?

$$\left(\frac{\text{furthest src}}{\text{closest tgt}}\right)^{k+1} = \left(\frac{\sqrt{2}}{3}\right)^{k+1} \quad \text{2D} \qquad \left(\frac{\sqrt{3}}{3}\right)^{k+1} \quad \text{3D}$$

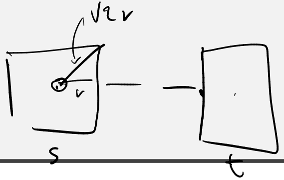

Q: Does this get better or worse as dimension increases?

Barnes-Hut (Single-Level): Computational Cost

What's the cost of this algorithm?

N # particles

K # terms in expansion

m # particles in a box

	How many	Cost	
Compute mpoles	N/m	Km	KN
Evaluate mpoles	N targets $\cdot \frac{N}{m}$	K	$\frac{N^2}{m}K$
9 close boxes	$9\frac{N}{m}$	m^2	$Nm \leftarrow$

$m = \sqrt{N}$

Barnes-Hut Single Level Cost: Observations

$$\text{cost} \sim \frac{N^2}{m} \sim N^{1.5}$$

Box Splitting

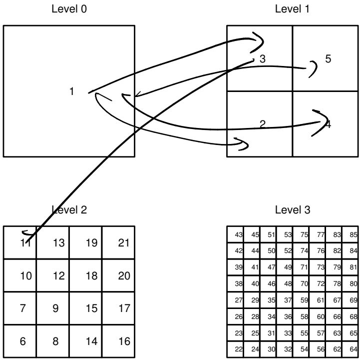

(Figure credit: G. Martinsson)

Level Count

How many levels?

?

Possible criterion:

m bounded

Box Sizes

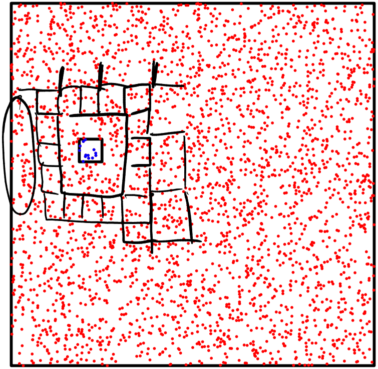

(Figure credit: G. Martinsson)

Want to evaluate all the source interactions with the targets in the box.

Q: What would be good sizes for source boxes? What's the requirement?

Multipole Sources

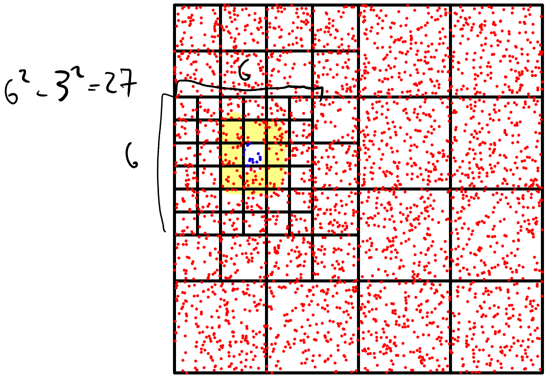

(Figure credit: G. Martinsson)

Data from which of these boxes could we bring in using multipole expansions? Does that depend on the type of expansion? (Taylor/special function vs skeletons)

Barnes-Hut: Constraints on Expansions

Data from which of these boxes could we bring in using multipole expansions? Does that depend on the type of expansion? (Taylor/special function vs skeletons)

Barnes-Hut: Multipole Accuracy

Idea:

- Don't use multipoles from the near neighbors (Instead: Compute interactions directly)
- Do use multipoles from non-near neighbors
- I.e. have a *buffer* of non-multipole source boxes around each target box

Note: Whether or not to use buffering is a *choice*, with the following tradeoff:

Pros (of buffering):

- Simple, constant-rank interactions
- Works for all expansion types

Cons:

- Trickier bookkeeping

Barnes-Hut: Box Properties

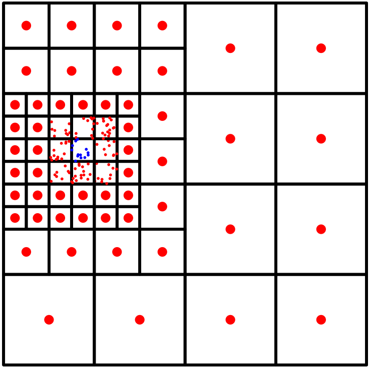

(Figure credit: G. Martinsson)

What properties do these boxes have?

Simple observation: The further, the bigger.

Barnes-Hut: Box Properties

Multipole Acceptance Crit (MAC)

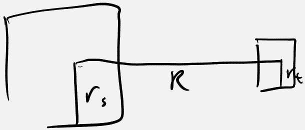

mpole accuracy $\approx \left(\frac{r_s \sqrt{2}}{R - r_t} \right)^{k+1}$

Barnes-Hut: Well-separated-ness

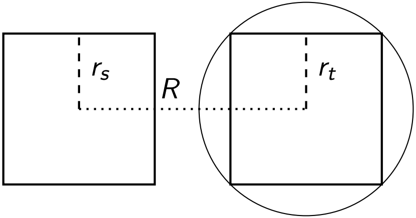

Convergent iff $r_s\sqrt{2} < R - r_t$. $(*)$
Convergent *if* (picture) $R \geqslant 3 \cdot \max(r_t, r_s)$ $(**)$
because $(*) \Leftrightarrow \left(r_t + \sqrt{2}r_s\right) < R$.

We'll make a new word for that: A pair of boxes satisfying the condition $(**)$ is called *well-separated*. Observations:

- This is just *one* choice. (the one we'll use anyway)
- One can play games here, based on a target accuracy. $\rightarrow$ *Multipole Acceptance Criterion* ('MAC') or *Admissibility Condition*

Barnes-Hut: Revised Cost Estimate

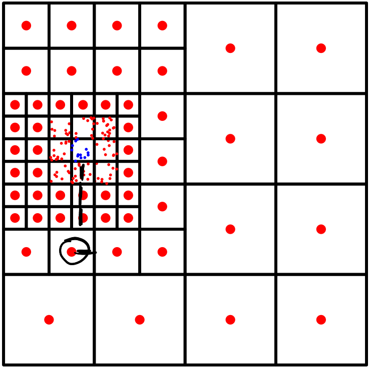

(Figure credit: G. Martinsson)

What is the cost of evaluating the target potentials, assuming that we know the multipole expansions already?

Barnes-Hut: Revised Cost Estimate

L # levels ~ log(N)?

N # particles

K # terms

m # particles in a box

Eval towards single tgt box

- 9 boxes of direct eval
- ≤ 27 src boxes per level

→ $O(L\,27\,K) = O(\log N)$

For all tgt. boxes is $O(N \log N)$

Barnes-Hut: Next Revised Cost Estimate

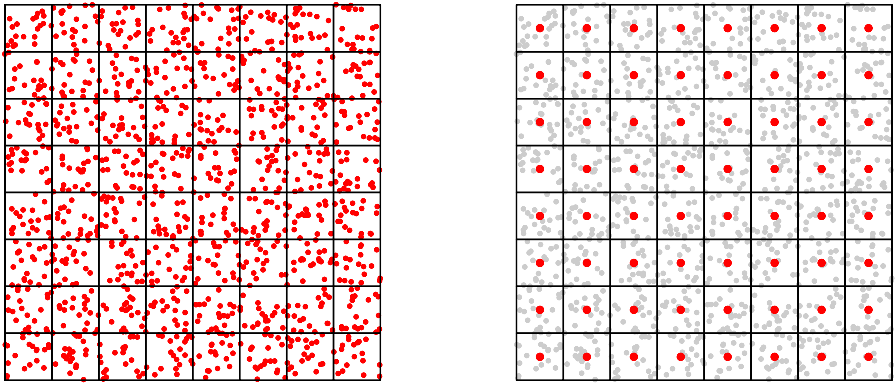

(Figure credit: G. Martinsson)

Summarize the algorithm (so far) and the associated cost.

Barnes-Hut: Next Revised Cost Estimate

Summarize the algorithm (so far) and the associated cost.

What	How many	Cost	Total
Compute mpoles	N_{src}	LK	$K N \log N$ Cold ↓ new $KN + K^2 N/m$
Evaluate mpoles	$N_{tgts} \cdot 27L$	K	$N K \log N$
9 close boxes	$9(N/m)$	m^2	N/m

Barnes-Hut: Putting Multipole Expansions to Work

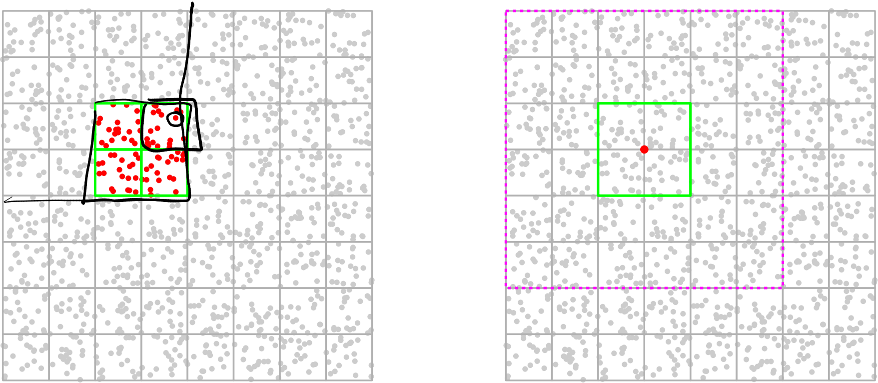

(Figure credit: G. Martinsson)

How could this process be sped up?

Barnes-Hut: Clumps of Boxes?

Observation: The amount of work does not really decrease as we go up the tree: Fewer boxes, but more particles in each of them.

But we already compute multipoles to summarize lower-level boxes...

Barnes-Hut: Putting Multipole Expansions to Work

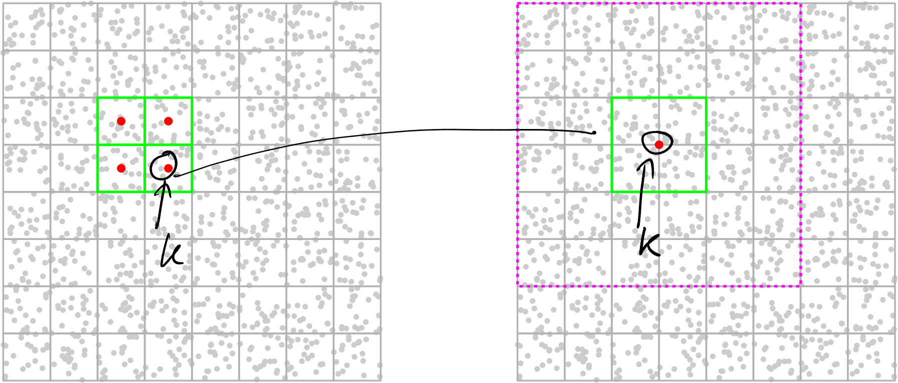

(Figure credit: G. Martinsson)

To get a new 'big' multipole from a 'small' multipole, we need a new mathematical tool.

Barnes-Hut: Translations

- shift center of mpole expansion to another center

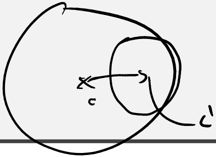

Cost of Multi-Level Barnes-Hut

Compute mpoles

Level	What	Cost	How Many
L (leaf)	Form mpoles	$m\,k$	N/m
L-1	mpole→mpole	k^2	(N/m)
L-2	— " —	k^2	$(N/m)/4$
L-3	— " —	k^2	$(N/m)/16$

$O(1) \cdot (N/m)$

Cost of Multi-Level Barnes-Hut: Observations

Observation: Multipole evaluation remains as the single most costly bit of this algorithm. *Fix?*

Idea: Exploit the tree structure also in performing this step.
If 'upward' translation of multipoles helped earlier, maybe 'downward' translation of *local* expansions can help now.

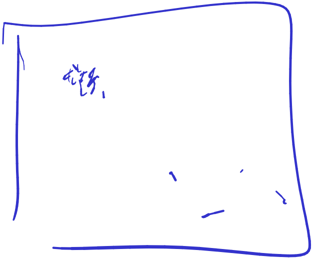

Using Multipole-to-Local

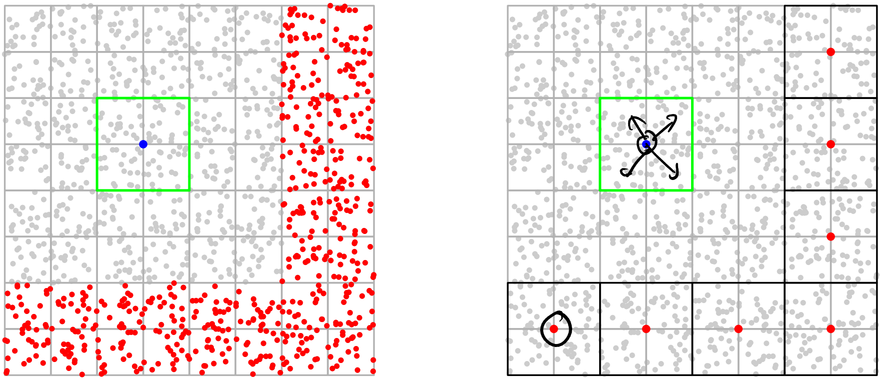

(Figure credit: G. Martinsson)

Come up with an algorithm that computes the interaction in the figure.

Using Multipole-to-Local

Come up with an algorithm that computes the interaction in the figure.

1. Form multipoles in each leaf
2. Translate multipole to local
3. Evaluate local.
4. Evaluate near neighbor interactions directly

Using Multipole-to-Local: Next Level

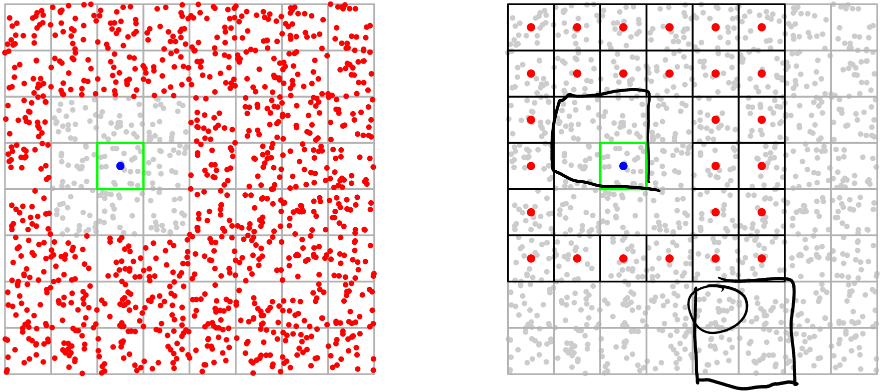

(Figure credit: G. Martinsson)

Assuming we retain information from the previous level, how can we obtain a valid local expansion on the target box?

Using Multipole-to-Local: Next Level

Assuming we retain information from the previous level, how can we obtain a valid local expansion on the target box?