

M397C

Fast algorithms in scientific computing – theory & practice

The Fast Multipole Method

April 14, 2022

Per-Gunnar Martinsson

The University of Texas at Austin

Problem definition: Consider the task of evaluating the sum

$$(1) \quad u_i = \sum_{j=1}^N G(\mathbf{x}_i, \mathbf{x}_j) q_j, \quad i = 1, 2, \dots, N,$$

where

$\{\mathbf{x}_i\}_{i=1}^N$ is a given set of points in a square Ω in the plane, where
 $\{q_i\}_{i=1}^N$ is a given set of real numbers which we call *sources*, and where
 $\{u_i\}_{i=1}^N$ is a sought set of real numbers which we call *potentials*.

The *kernel* G is given by

$$(2) \quad G(\mathbf{x}, \mathbf{y}) = \begin{cases} \log(\mathbf{x} - \mathbf{y}), & \text{when } \mathbf{x} \neq \mathbf{y} \\ 0 & \text{when } \mathbf{x} = \mathbf{y}. \end{cases}$$

Recall: A point $\mathbf{x} \in \mathbb{R}^2$ is represented by the complex number

$$\mathbf{x} = x_1 + i x_2 \in \mathbb{C}.$$

Then the kernel in (2) is a complex representation of the fundamental solution to the Laplace equation in \mathbb{R}^2 since

$$\log |\mathbf{x} - \mathbf{y}| = \text{Real}(\log(\mathbf{x} - \mathbf{y})).$$

(The factor of $-1/2\pi$ is suppressed.)

Special case: Sources and targets are separate


Charges q_j at locations $\{\mathbf{y}_j\}_{j=1}^N$.

Potentials u_i at locations $\{\mathbf{x}_i\}_{i=1}^M$.

$$u_i = \sum_{j=1}^N \log(\|\mathbf{x}_i - \mathbf{y}_j\|) q_j, \quad i = 1, 2, \dots, M$$

$$u_i = u(\mathbf{x}_i), \quad i = 1, 2, \dots, M$$

$$u(\mathbf{x}) = \sum_{j=1}^N \log(\|\mathbf{x} - \mathbf{y}_j\|) q_j$$



Direct evaluation

Cost is $O(MN)$.

Special case: Sources and targets are separate


Charges q_j at locations $\{\mathbf{y}_j\}_{j=1}^N$.

Potentials u_i at locations $\{\mathbf{x}_i\}_{i=1}^M$.

$$u_i = \sum_{j=1}^N \log(\mathbf{x}_i - \mathbf{y}_j) q_j, \quad i = 1, 2, \dots, M$$

$$u_i = u(\mathbf{x}_i), \quad i = 1, 2, \dots, M$$

$$u(\mathbf{x}) = \sum_{j=1}^N \log(\mathbf{x} - \mathbf{y}_j) q_j$$



Direct evaluation

Cost is $O(MN)$.

But recall that $\log(\mathbf{x} - \mathbf{y})$ admits the *separation of variables*

$$\log(\mathbf{x} - \mathbf{y}) = \log(\mathbf{x} - \mathbf{c}) \mathbf{1} + \sum_{\rho=1}^{\infty} \frac{-1}{\rho} \frac{1}{(\mathbf{x} - \mathbf{c})^{\rho}} (\mathbf{y} - \mathbf{c})^{\rho}.$$

(Recall that \mathbf{c} is the center of the source box.)

Special case: Sources and targets are separate


Charges q_j at locations $\{\mathbf{y}_j\}_{j=1}^N$.

Potentials u_i at locations $\{\mathbf{x}_i\}_{i=1}^M$.

$$u_i = \sum_{j=1}^N \log(\mathbf{x}_i - \mathbf{y}_j) q_j, \quad i = 1, 2, \dots, M$$

$$u_i = u(\mathbf{x}_i), \quad i = 1, 2, \dots, M$$

$$u(\mathbf{x}) = \sum_{j=1}^N \log(\mathbf{x} - \mathbf{y}_j) q_j$$



Multipole expansion:

It follows that $u(\mathbf{x}) = \log(\mathbf{x} - \mathbf{c}) \hat{q}_0 + \sum_{p=1}^{\infty} \frac{1}{(\mathbf{x} - \mathbf{c})^p} \hat{q}_p$.

where $\hat{q}_0 = \sum_j^n q_j$ and $\hat{q}_p = -\frac{1}{p} \sum_{j=1}^N (\mathbf{y}_j - \mathbf{c})^p q_j$.

Special case: Sources and targets are separate

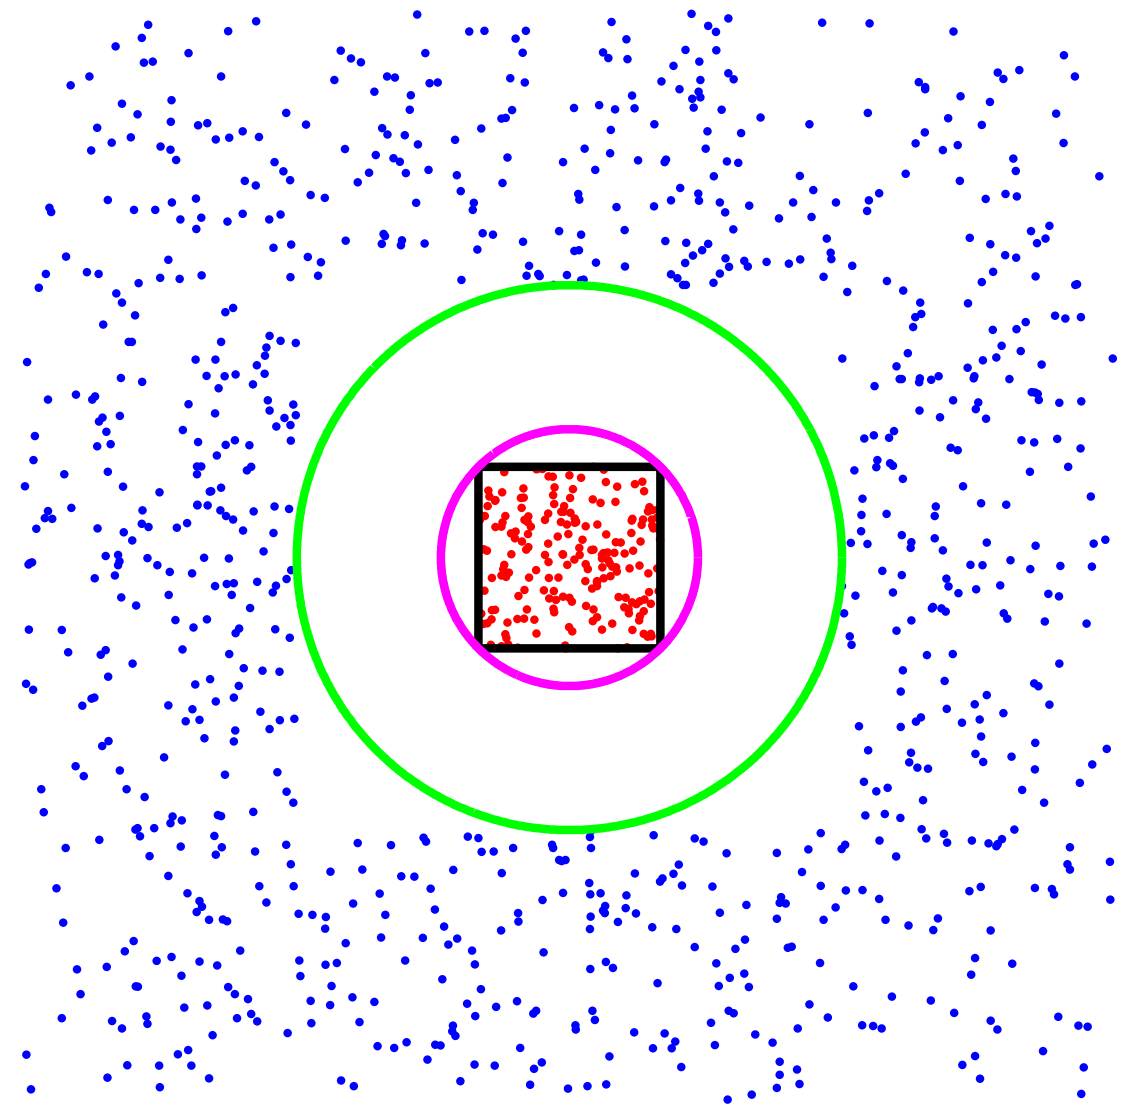
Charges q_j at locations $\{\mathbf{y}_j\}_{j=1}^N$.

Potentials u_i at locations $\{\mathbf{x}_i\}_{i=1}^M$.

$$u_i = \sum_{j=1}^N \log(\mathbf{x}_i - \mathbf{y}_j) q_j, \quad i = 1, 2, \dots, M$$

$$u_i = u(\mathbf{x}_i), \quad i = 1, 2, \dots, M$$

$$u(\mathbf{x}) = \sum_{j=1}^N \log(\mathbf{x} - \mathbf{y}_j) q_j$$



Multipole expansion — truncated to $P + 1$ terms:

It follows that $u(\mathbf{x}) = \log(\mathbf{x} - \mathbf{c})\hat{q}_0 + \sum_{p=1}^P \frac{1}{(\mathbf{x} - \mathbf{c})^p} \hat{q}_p + E_P$ where $\hat{q}_p = -\frac{1}{p} \sum_{j=1}^N (\mathbf{y}_j - \mathbf{c})^p q_j$.

The approximation error E_P scales as $E_P \sim \left(\frac{r}{R}\right)^P = \left(\frac{\sqrt{2}a}{3a}\right)^P = \left(\frac{\sqrt{2}}{3}\right)^P$, where

$r = \sqrt{2}a$ is the radius of the magenta circle, and $R = 3a$ is the radius of the green circle.

Special case: Sources and targets are separate


Charges q_j at locations $\{\mathbf{y}_j\}_{j=1}^N$.

Potentials u_i at locations $\{\mathbf{x}_i\}_{i=1}^M$.

$$u_i = \sum_{j=1}^N \log(\|\mathbf{x}_i - \mathbf{y}_j\|) q_j, \quad i = 1, 2, \dots, M$$

$$u_i = u(\mathbf{x}_i), \quad i = 1, 2, \dots, M$$

$$u(\mathbf{x}) = \sum_{j=1}^N \log(\|\mathbf{x} - \mathbf{y}_j\|) q_j$$



Multipole expansion — truncated to $P + 1$ terms — costs:

Evaluate $\hat{q}_p = -\frac{1}{p} \sum_{j=1}^N (\mathbf{y}_j - \mathbf{c})^p q_j$ for $p = 0, 1, \dots, P$ — *cost is $O(NP)$.*

Evaluate $u_i = \log(\|\mathbf{x}_i - \mathbf{c}\|) \hat{q}_0 + \sum_{p=1}^P \frac{1}{(\|\mathbf{x}_i - \mathbf{c}\|)^p} \hat{q}_p$ — *cost is $O(MP)$.*

The cost has been reduced from $O(MN)$ to $O(P(M + N))$.

You can view the multipole expansion as a *matrix factorization*.

Let \mathbf{A} denote the $m \times n$ matrix with entries

$$\mathbf{A}(i, j) = \log(\mathbf{x}_i - \mathbf{y}_j).$$

Then


$$\begin{array}{ccc} \mathbf{A} & \approx & \tilde{\mathbf{B}} \quad \mathbf{C} \\ m \times n & & m \times (P + 1) \quad (P + 1) \times n \end{array}$$

where

$$\mathbf{C} \text{ is the } (P + 1) \times n \text{ matrix with entries } \mathbf{C}(p, j) = \begin{cases} 1 & p = 0 \\ -(1/p) (\mathbf{y}_j - \mathbf{c})^p & p \neq 0 \end{cases}$$


$$\tilde{\mathbf{B}} \text{ is the } m \times (P + 1) \text{ matrix with entries } \tilde{\mathbf{B}}_{ip} = \begin{cases} \log(\mathbf{x}_i - \mathbf{c}) & p = 0 \\ 1/(\mathbf{x}_i - \mathbf{c})^p & p \neq 0 \end{cases}$$

You could also view this as a commutative diagram



Definition: Let Ω be a square with center $\mathbf{c} = (c_1, c_2)$ and side length $2a$. Then we say that a point $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$ is *well-separated* from Ω if

$$\max(|x_1 - c_1|, |x_2 - c_2|) \geq 3a.$$



Any point on or outside of the dashed square is well-separated from Ω .

Blue points are well-separated. Red points are not.

Definition:

Let Ω be a square with center \mathbf{c} containing sources $\{q_j\}_{j=1}^n$ at locations $\{\mathbf{y}_j\}_{j=1}^n$.

The *outgoing expansion* of Ω (to order P) is the vector $\hat{\mathbf{q}} \in \mathbb{C}^{P+1}$ with numbers


$$\hat{q}_0 = \sum_{j=1}^n q_j,$$
$$\hat{q}_p = -\frac{1}{p} \sum_{j=1}^n (\mathbf{y}_j - \mathbf{c})^j q_j, \quad p = 1, 2, 3, \dots, P.$$

The outgoing expansion compactly encodes source strengths and source locations.

It allows you to evaluate the potential u caused by the sources in Ω (to precision that depends on P and the distance to the sources).

“outgoing expansion” = “multipole expansion”

Single-level Barnes-Hut




We seek to evaluate all pairwise interactions between N particles in a box; for now, assume that the particle locations $\{\mathbf{x}_i\}_{i=1}^N$ are fairly evenly distributed.

We cast the problem as a matrix-vector multiplication $\mathbf{u} = \mathbf{A}\mathbf{q}$, where the $N \times N$ matrix \mathbf{A} has entries

$$\mathbf{A}(i, j) = \log(\mathbf{x}_i - \mathbf{x}_j).$$

Single-level Barnes-Hut




Place a square grid of boxes on top of the computational box.

Assume each box holds about m particles (so there are about N/m boxes).

Given a tolerance ε , pick P so that, roughly, $(\sqrt{2}/3)^P < \varepsilon$ (... details left out ...).


For each box, compute its outgoing expansion.

Single-level Barnes-Hut



How do you evaluate the potentials at the blue locations?

Single-level Barnes-Hut



How do you evaluate the potentials at the blue locations?

Directly evaluate interactions with the particles close by.

For all long-distance interactions, use the outgoing expansion!

Cost of the single-level Barnes-Hut algorithm:

Let m denote the number of particles in a box.

For each particle, we need to do the following:

Step 1: Evaluate the P outgoing moments: P

Step 2: Evaluate potentials from outgoing expansions: $((N/m) - 9) P$

Step 3: Evaluate potentials from close particles: $9 m$

Using that P is a smallish constant we find

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

Set $m \sim N^{1/2}$ to obtain:

$$\text{cost} \sim N^{1.5}.$$

We're doing better than $O(N^2)$ but still not great.

Notation — single level Barnes-Hut:

Partition the box Ω into smaller boxes and label them:

8	16	24	32	40	48	56	64
7	15	23	31	39	47	55	63
6	14	22	30	38	46	54	62
5	13	21	29	37	45	53	61
4	12	20	28	36	44	52	60
3	11	19	27	35	43	51	59
2	10	18	26	34	42	50	58
1	9	17	25	33	41	49	57

For a box τ , let $\mathcal{L}_{\tau}^{(\text{nei})}$ denote its neighbors, and $\mathcal{L}_{\tau}^{(\text{far})}$ denote the remaining boxes.

Notation — single level Barnes-Hut:

Partition the box Ω into smaller boxes and label them:

8	16	24	32	40	48	56	64
7	15	23	31	39	47	55	63
6	14	22	30	38	46	54	62
5	13	21	29	37	45	53	61
4	12	20	28	36	44	52	60
3	11	19	27	35	43	51	59
2	10	18	26	34	42	50	58
1	9	17	25	33	41	49	57

For a box τ , let $\mathcal{L}_{\tau}^{(\text{nei})}$ denote its neighbors, and $\mathcal{L}_{\tau}^{(\text{far})}$ denote the remaining boxes.

The box $\tau = 21$ is marked with red.

Notation — single level Barnes-Hut:

Partition the box Ω into smaller boxes and label them:

8	16	24	32	40	48	56	64
7	15	23	31	39	47	55	63
6	14	22	30	38	46	54	62
5	13	21	29	37	45	53	61
4	12	20	28	36	44	52	60
3	11	19	27	35	43	51	59
2	10	18	26	34	42	50	58
1	9	17	25	33	41	49	57

For a box τ , let $\mathcal{L}_{\tau}^{(\text{nei})}$ denote its neighbors, and $\mathcal{L}_{\tau}^{(\text{far})}$ denote the remaining boxes.

The box $\tau = 21$ is marked with red.

$\mathcal{L}_{21}^{(\text{nei})} = \{12, 13, 14, 20, 22, 28, 29, 30\}$ are the blue boxes.

Notation — single level Barnes-Hut:

Partition the box Ω into smaller boxes and label them:

8	16	24	32	40	48	56	64
7	15	23	31	39	47	55	63
6	14	22	30	38	46	54	62
5	13	21	29	37	45	53	61
4	12	20	28	36	44	52	60
3	11	19	27	35	43	51	59
2	10	18	26	34	42	50	58
1	9	17	25	33	41	49	57

For a box τ , let $\mathcal{L}_{\tau}^{(\text{nei})}$ denote its neighbors, and $\mathcal{L}_{\tau}^{(\text{far})}$ denote the remaining boxes.

The box $\tau = 21$ is marked with red.

$\mathcal{L}_{21}^{(\text{nei})} = \{12, 13, 14, 20, 22, 28, 29, 30\}$ are the blue boxes.

$\mathcal{L}_{21}^{(\text{far})} = \dots$ are the green boxes.

Notation — single level Barnes-Hut:

Partition the box Ω into smaller boxes and label them:

8	16	24	32	40	48	56	64
7	15	23	31	39	47	55	63
6	14	22	30	38	46	54	62
5	13	21	29	37	45	53	61
4	12	20	28	36	44	52	60
3	11	19	27	35	43	51	59
2	10	18	26	34	42	50	58
1	9	17	25	33	41	49	57

For a box τ , let $\mathcal{L}_\tau^{(\text{nei})}$ denote its neighbors, and $\mathcal{L}_\tau^{(\text{far})}$ denote the remaining boxes.

The box $\tau = 21$ is marked with red.

$\mathcal{L}_{21}^{(\text{nei})} = \{12, 13, 14, 20, 22, 28, 29, 30\}$ are the blue boxes.

$\mathcal{L}_{21}^{(\text{far})} = \dots$ are the green boxes.

Let J_τ denote an index vector marking which particles belong to τ :

$$j \in J_\tau \quad \Leftrightarrow \quad \mathbf{x}_j \text{ is in box } \tau.$$

Single-level Barnes-Hut

Compute the outgoing expansions on all boxes:

loop over all boxes τ

$$\hat{\mathbf{q}}_{\tau} = \mathbf{C}_{\tau} \mathbf{q}(J_{\tau})$$

end loop

Evaluate the far field potentials.

Each box τ aggregates the contributions from all well-separated boxes:

$\mathbf{u} = \mathbf{0}$

loop over all boxes τ

loop over all $\sigma \in \mathcal{L}_{\tau}^{(\text{far})}$ (i.e. all σ that are well-separated from τ)

$$\mathbf{u}(J_{\tau}) = \mathbf{u}(J_{\tau}) + \tilde{\mathbf{B}}_{\tau,\sigma} \hat{\mathbf{q}}_{\sigma}$$

end loop

end loop

Evaluate the near field interactions:


loop over all leaf boxes τ

$$\mathbf{u}(J_{\tau}) = \mathbf{u}(J_{\tau}) + \mathbf{A}(J_{\tau}, J_{\tau}) \mathbf{q}(J_{\tau}) + \sum_{\sigma \in \mathcal{L}_{\tau}^{(\text{nei})}} \mathbf{A}(J_{\tau}, J_{\sigma}) \mathbf{q}(J_{\sigma})$$


end

To get the asymptotic cost down further, we need a *hierarchy of boxes* (or a “tree of boxes”) on the computational domain:


Level 0




Level 1




Level 2




Level 3




For each box in the tree, compute its outgoing expansion.



How do you find the potential at the locations marked in blue?



Tessellate the remainder of the domain using as large boxes as you can, with the constraint that *the target box has to be well-separated from every box that is used.*



Then replace the original sources in each well-separated box by the corresponding multipole expansion.

The work required to evaluate one potential is now $O(\log N)$.

Cost of the multi-level Barnes-Hut algorithm:

Suppose there are L levels in the tree, and that there are about m particles in each box so that $L \approx \log_4(N/m)$.

We let m be a fixed number (say $m \approx 200$) so $L \sim \log(N)$.

Observation: On each level, there are at most 27 well-separated boxes.

For each particle \mathbf{x}_j , we need to do the following:

Step 1: Evaluate the P outgoing moments for the L boxes holding \mathbf{x}_j :	$L P$
Step 2: Evaluate potentials from outgoing expansions:	$27 L P$
Step 3: Evaluate potentials from neighbors:	$9 m$

Using that P is a smallish constant (say $P = 20$) we find

$$\text{cost} \sim N L \sim N \log(N).$$

This is not bad.

Multi-level Barnes-Hut

Compute the outgoing expansions on all boxes:

loop over all boxes τ on all levels

$$\hat{\mathbf{q}}_{\tau} = \mathbf{C}_{\tau} \mathbf{q}(J_{\tau})$$

end loop

For each box τ tessellate Ω into a minimal collection of boxes from which τ is well-separated, and evaluate the far-field potential:

$$\mathbf{u} = 0$$

loop over all boxes τ

loop over all $\sigma \in \mathcal{L}_{\tau}^{(\text{BH})}$ (i.e. all σ that are well-separated from τ)

$$\mathbf{u}(J_{\tau}) = \mathbf{u}(J_{\tau}) + \tilde{\mathbf{B}}_{\tau,\sigma} \hat{\mathbf{q}}_{\sigma}$$

end loop

end loop

Evaluate the near field interactions:

loop over all leaf boxes τ

$$\mathbf{u}(J_{\tau}) = \mathbf{u}(J_{\tau}) + \mathbf{A}(J_{\tau}, J_{\tau}) \mathbf{q}(J_{\tau}) + \sum_{\sigma \in \mathcal{L}_{\tau}^{(\text{nei})}} \mathbf{A}(J_{\tau}, J_{\sigma}) \mathbf{q}(J_{\sigma})$$

end



The Barnes-Hut algorithm has asymptotic cost $O(N \log N)$ (at constant precision).

We will next improve the accuracy to the optimal $O(N)$.

Recall: We partition into $L \sim \log N$ levels. Then the costs are:

	<i>Asymptotic cost</i>	<i>How get to $O(N)$?</i>
Stage 1 — compute all outgoing expansions: Each particle communicates with all L boxes that contain it.	$O(N \log N)$	Easy!
Stage 2 — compute all far field potentials: Each particle gathers contributions from $\sim 27L$ outgoing expansions.	$O(N \log N)$???
Stage 3 — compute all near field potentials: Each particle gathers contributions directly from all particles in the same box, and from all particles in the 8 immediate neighbors.	$O(N)$	No need to fix!

Reducing the cost of computing all outgoing expansions from $O(N \log N)$ to $O(N)$:





For every leaf box τ , we directly compute the outgoing expansion from the source vector

$$\hat{\mathbf{q}}_{\tau} = \mathbf{C}_{\tau} \mathbf{q}(J_{\tau}).$$

(Just as before.)

Reducing the cost of computing all outgoing expansions from $O(N \log N)$ to $O(N)$:



Now consider a box Ω_τ made up of four leaves: $\Omega_\tau = \Omega_{\sigma_1} \cup \Omega_{\sigma_2} \cup \Omega_{\sigma_3} \cup \Omega_{\sigma_4}$



We seek an outgoing expansion that is valid outside the dotted magenta line.

In this region, the outgoing expansions of the children $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$ are valid.

“Move” these expansions via a so called *outgoing-from-outgoing translation operator*:

$$\hat{\mathbf{q}}_\tau = \sum_{j=1}^4 \mathbf{T}_{\tau, \sigma_j}^{(\text{ofO})} \hat{\mathbf{q}}_{\sigma_j}.$$

Reducing the cost of computing all outgoing expansions from $O(N \log N)$ to $O(N)$:



Now consider a box Ω_τ made up of four leaves: $\Omega_\tau = \Omega_{\sigma_1} \cup \Omega_{\sigma_2} \cup \Omega_{\sigma_3} \cup \Omega_{\sigma_4}$

We seek an outgoing expansion that is valid outside the dotted magenta line.

In this region, the outgoing expansions of the children $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$ are valid.

“Move” these expansions via a so called *outgoing-from-outgoing translation operator*:

$$\hat{\mathbf{q}}_\tau = \sum_{j=1}^4 \mathbf{T}_{\tau, \sigma_j}^{(\text{of})} \hat{\mathbf{q}}_{\sigma_j}.$$

The *outgoing-from-outgoing* translation operator $\mathbf{T}_{\tau,\sigma}^{(\text{ofo})}$


Let τ be a box (green).

Let \mathbf{c}_τ be the center of τ (black).

Let σ denote a box contained in τ .

Let \mathbf{c}_σ denote the center of σ (red).

Let $\hat{\mathbf{q}}_\sigma$ be outgoing expansion of σ .



$\mathbf{T}_{\tau,\sigma}^{(\text{ofo})}$ constructs the outgoing expansion of τ from the outgoing expansion of σ

$$\hat{\mathbf{q}}_\tau = \mathbf{T}_{\tau,\sigma}^{(\text{ofo})} \hat{\mathbf{q}}_\sigma$$

$(P+1) \times 1 \quad (P+1) \times (P+1) \quad (P+1) \times 1$

With $\mathbf{d} = \mathbf{c}_\sigma - \mathbf{c}_\tau$, $\mathbf{T}_{\tau,\sigma}^{(\text{ofo})}$ is a lower tridiagonal matrix with entries

$$\mathbf{T}_{\tau,\sigma,0,0}^{(\text{ofo})} = 1$$

$$\mathbf{T}_{\tau,\sigma,p,0}^{(\text{ofo})} = -\frac{1}{p} \mathbf{d} \quad 1 \leq p \leq P$$

$$\mathbf{T}_{\tau,\sigma,p,q}^{(\text{ofo})} = \binom{p}{q} \mathbf{d}^{p-q} \quad 1 \leq q \leq p \leq P.$$

Cost of constructing all outgoing expansions:

(Recall: $L \sim \log N$ is the number of levels. P is the length of the expansion.)

Level L (the leaves)	“outgoing from sources”	NP
Level $L - 1$ (next coarser)	“outgoing from outgoing”	$(N_{\text{boxes}}) P^2$
Level $L - 2$	“outgoing from outgoing”	$(N_{\text{boxes}}/4) P^2$
Level $L - 3$	“outgoing from outgoing”	$(N_{\text{boxes}}/16) P^2$
Level $L - 4$	“outgoing from outgoing”	$(N_{\text{boxes}}/64) P^2$
⋮	⋮	⋮
Total:		$NP + N_{\text{boxes}} P^2$

We succeeded in attaining $O(N)$ cost for computing all outgoing expansions.

Next we will devise an $O(N)$ scheme for computing so called *incoming expansions*.


The incoming expansion

Let Ω_τ be a box (green).

Let \mathbf{c}_τ be the center of τ (black).

Let $I_\tau^{(\text{far})}$ denote a list of all sources well-separated from τ (red), and let ϕ denote the potential

$$\phi(\mathbf{x}) = \sum_{j \in I_\tau^{(\text{far})}} \log(\mathbf{x}_j - \mathbf{y}_j) q_j, \quad \mathbf{x} \in \Omega_\tau.$$



The *incoming expansion* of τ is a vector $\hat{\mathbf{u}} = [\hat{u}_p]_{p=0}^P$ of complex numbers such that

$$(3) \quad \phi(\mathbf{x}) \approx \sum_{p=0}^P \hat{u}_p (\mathbf{x} - \mathbf{c}_\tau)^p, \quad \mathbf{x} \in \Omega_\tau.$$

The incoming expansion is a compact representation of the field generated by sources that are well-separated from Ω_τ (it encodes both the source locations and magnitudes).

Recall:

$$\phi(\mathbf{x}) \approx \sum_{p=0}^P \hat{u}_p (\mathbf{x} - \mathbf{c}_\tau)^p, \quad \mathbf{x} \in \Omega_\tau.$$

Recall that each $\hat{u}_p = \hat{u}_p^r + i \hat{u}_p^i$ is a complex number.

Taking real parts, and using polar coordinates,



$$\mathbf{x} - \mathbf{c}_\tau = r e^{i\theta},$$

we get

$$\begin{aligned} \text{Real}(\phi(\mathbf{x})) = & \hat{u}_0^r + \\ & \hat{u}_1^r r^1 \cos(1\theta) - \hat{u}_1^i r^1 \sin(1\theta) + \dots \\ & \hat{u}_2^r r^2 \cos(2\theta) - \hat{u}_2^i r^2 \sin(2\theta) + \dots \\ & \hat{u}_3^r r^3 \cos(3\theta) - \hat{u}_3^i r^3 \sin(3\theta) + \dots \end{aligned}$$

Classical expansion in harmonics.

Computing the incoming expansions for all boxes in $O(N)$ operations





We seek to construct the incoming expansion for box τ (marked in green).

We use the outgoing expansions for all well-separated boxes:

$$\hat{\mathbf{u}}_{\tau} = \sum_{\sigma \in \mathcal{L}_{\tau}^{(\text{int})}} \mathbf{T}_{\tau, \sigma}^{(\text{ifo})} \hat{\mathbf{q}}_{\sigma}$$



where $\mathbf{T}_{\tau, \sigma}$ is the *incoming-from-outgoing* translation operator, and $\mathcal{L}_{\tau}^{(\text{int})}$ is the *interaction list* of box τ .

Computing the incoming expansions for all boxes in $O(N)$ operations



We seek to construct the incoming expansion for box τ (marked in green):

Computing the incoming expansions for all boxes in $O(N)$ operations



We seek to construct the incoming expansion for box τ (marked in green):

Transfer the incoming expansion from the parent box, ν , and then add all contributions from boxes in the interaction list:

$$\hat{\mathbf{u}}_{\tau} = \mathbf{T}_{\tau, \nu}^{(\text{ifi})} \hat{\mathbf{u}}_{\nu} + \sum_{\sigma \in \mathcal{L}_{\tau}^{(\text{int})}} \mathbf{T}_{\tau, \sigma}^{(\text{ifo})} \hat{\mathbf{q}}_{\sigma}.$$

Definition of “interaction list”:

For a box τ , define its *interaction list* $\mathcal{L}_\tau^{(\text{int})}$ as the set of all boxes σ such that:

1. σ and τ populate the same level of the tree.
2. σ and τ are well-separated.
3. The *parents* of σ and τ touch.


The *incoming-from-outgoing* translation operator $\mathbf{T}_{\tau,\sigma}^{(\text{ifo})}$

Let σ be a source box (red) with center \mathbf{c}_σ .

Let τ be a target box (blue) with center \mathbf{c}_τ .

Let $\hat{\mathbf{q}}_\sigma$ be the outgoing expansion of σ .

Let $\hat{\mathbf{u}}_\tau$ represent the potential in τ caused by sources in σ .




$\mathbf{T}_{\tau,\sigma}^{(\text{ifo})}$ constructs the incoming expansion of τ from the outgoing expansions of σ :

$$\begin{array}{ccccc} \hat{\mathbf{u}}_\tau & = & \mathbf{T}_{\tau,\sigma}^{(\text{ifo})} & & \hat{\mathbf{q}}_\sigma \\ (P+1) \times 1 & & (P+1) \times (P+1) & & (P+1) \times 1 \end{array}$$

With $\mathbf{d} = \mathbf{c}_\sigma - \mathbf{c}_\tau$, $\mathbf{T}_{\tau,\sigma}^{(\text{ifo})}$ is a matrix with entries

$$\mathbf{T}_{\tau,\sigma,p,q}^{(\text{ifo})} = ?$$

Computing the incoming expansions for all boxes in $O(N)$ operations



Finally, we construct the potential for a leaf box τ .

Far-field contributions are evaluated by simply expanding the incoming expansion.

Near-field contributions are evaluated directly (uncompressed):

$$\mathbf{u}(l_\tau) = \mathbf{T}_{\tau,\nu}^{(\text{tfi})} \hat{\mathbf{u}}_\tau + \mathbf{A}(l_\tau, l_\tau) \mathbf{q}(l_\tau) + \sum_{\sigma \in \mathcal{L}_\tau^{(\text{near})}} \mathbf{A}(l_\tau, l_\sigma) \mathbf{q}(l_\sigma)$$

THE CLASSICAL FAST MULTIPOLE METHOD IN \mathbb{R}^2

1. Construct the tree and all “interaction lists.”
2. For each leaf node, compute its outgoing expansion directly from the sources in the box via the *outgoing-from-sources operator*.
3. For each parent node, compute its outgoing expansion by merging the expansions of its children via the *outgoing-from-outgoing operator*.
4. For each node, compute its incoming expansion by transferring the incoming expansion of its parent (via the *incoming-from-incoming operator*), and then add the contributions from all sources in its interaction list (via the *incoming-from-outgoing operator*).
5. For each leaf node, evaluate the incoming expansion at the targets (via the *targets-from-incoming operator*), and compute near-field interactions directly.

Set $\hat{\mathbf{u}}_\tau = \mathbf{0}$ and $\hat{\mathbf{q}}_\tau = \mathbf{0}$ for all τ .

loop over all leaf nodes τ

$$\hat{\mathbf{q}}_\tau = \mathbf{T}_\tau^{(\text{ofs})} \mathbf{q}(J_\tau)$$

end loop

loop over levels $\ell = L, L - 1, \dots, 2$

loop over all nodes τ on level ℓ

$$\hat{\mathbf{q}}_\tau = \sum_{\sigma \in \mathcal{L}_\tau^{(\text{child})}} \mathbf{T}_{\tau,\sigma}^{(\text{ofc})} \hat{\mathbf{q}}_\sigma$$

end loop

end loop

loop over all nodes τ

$$\hat{\mathbf{u}}_\tau = \hat{\mathbf{u}}_\tau + \sum_{\sigma \in \mathcal{L}_\tau^{(\text{int})}} \mathbf{T}_{\tau,\sigma}^{(\text{ifo})} \hat{\mathbf{q}}_\sigma.$$

end loop

loop over levels $\ell = 2, 3, 4, \dots, L - 1$

loop over all nodes τ on level ℓ

loop over all children σ of τ

$$\hat{\mathbf{u}}_\sigma = \hat{\mathbf{u}}_\sigma + \mathbf{T}_{\sigma,\tau}^{(\text{ifi})} \hat{\mathbf{u}}_\tau.$$

end loop

end loop

end loop

loop over all leaf nodes τ

$$\mathbf{u}(J_\tau) = \mathbf{T}_\tau^{(\text{tfi})} \hat{\mathbf{u}}_\tau$$

end loop

loop over all leaf nodes τ

$$\begin{aligned} \mathbf{u}(J_\tau) &= \mathbf{u}(J_\tau) + \mathbf{A}(J_\tau, J_\tau) \mathbf{q}(J_\tau) \\ &\quad + \sum_{\sigma \in \mathcal{L}_\tau^{(\text{nei})}} \mathbf{A}(J_\tau, J_\sigma) \mathbf{q}(J_\sigma) \end{aligned}$$

end loop

Important: This introduction to the FMM was brief and *highly* incomplete!

Non-uniform particle distributions: These require the use of *adaptive trees*.

Everything generalizes nicely, but two additional “lists” and “translation operators” are needed. *The ability to handle non-uniform distributions is a key advantage of the FMM!*

Extension to 3D: In 2D, the “interaction list” has at most 27 elements ($27 = 6^2 - 3^2$). In 3D, it typically has *189 elements* ($189 = 6^3 - 3^3$). Moreover, multipole expansions in 3D converge much more slowly. Instead of $\varepsilon \approx (\sqrt{3}/2)^{p/2}$, we have $\varepsilon \approx (1/\sqrt{3})^{\sqrt{p}}$. This makes the cost of evaluating all “incoming-from-outgoing” translation operators expensive. This cost can be greatly reduced, however, by using more sophisticated translation operators (so called “diagonal forms”).

Helmholtz equation: If the whole computational domain is small in terms of wave-lengths, then the FMM for Helmholtz equation looks very similar to what we just showed. However, if the box is large in terms of wave-lengths, then very different machinery is required. The high-frequency FMM is much more subtle! It does not rely on rank-considerations alone.

Coding the FMM well is not that easy: If you can, it is recommended to use a packaged version. The FMM for Laplace’s equation in two dimensions is OK, but 3D is harder, and wideband Helmholtz is substantially harder.