The Fast Multipole Method

and other Fast Summation Techniques

Gunnar Martinsson
The University of Colorado at Boulder

Problem definition: Consider the task of evaluation the sum

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

where

$\{\boldsymbol{x}_i\}_{i=1}^N$ 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 *charges,* and where

$\{u_i\}_{i=1}^N$ is a sought set of real numbers which we call *potentials.*

The *kernel* G is given by

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

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

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

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

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

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

Special case: Sources and targets are separate

Charges q_j at locations $\{y_j\}_{j=1}^{N}$.

Potentials u_i at locations $\{x_i\}_{i=1}^{M}$.

$$u_i = \sum_{j=1}^{N} q_j \log(x_i - y_j), \qquad i = 1,\, 2,\, \dots,\, M$$

$$u_i = u(x_i), \qquad i = 1,\, 2,\, \dots,\, M$$

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

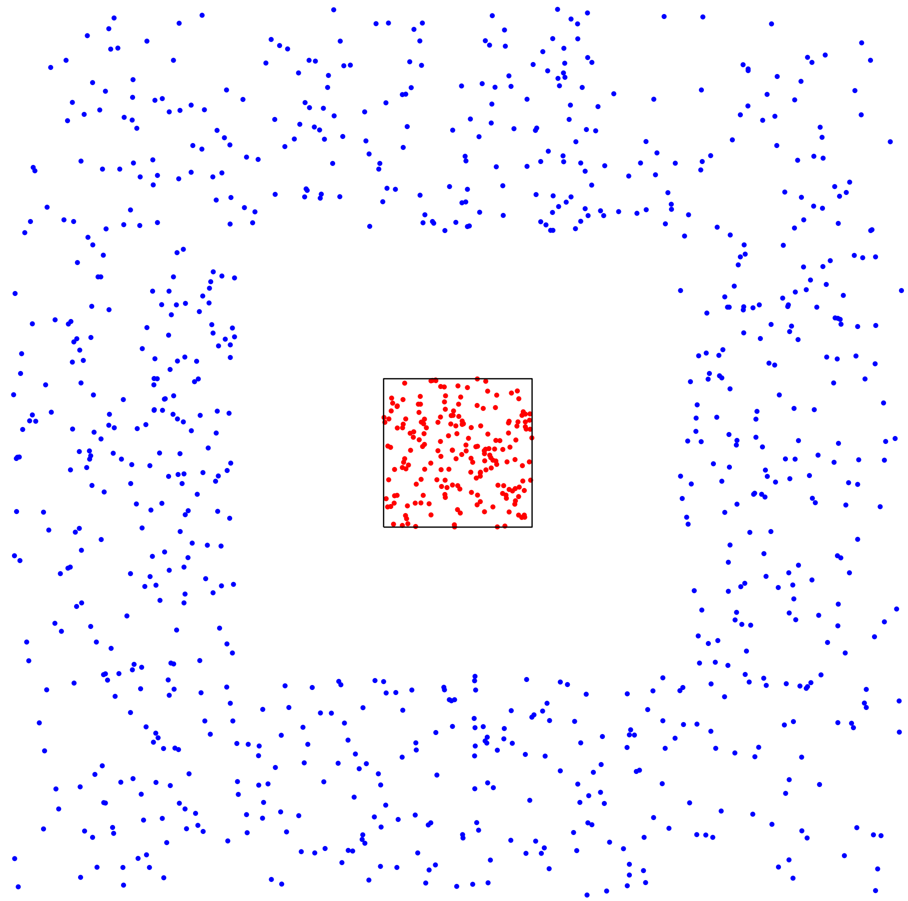

Direct evaluation

Cost is $O(M\,N)$.

Special case: Sources and targets are separate

Charges q_j at locations $\{y_j\}_{j=1}^{N}$.

Potentials u_i at locations $\{x_i\}_{i=1}^{M}$.

$$u_i = \sum_{j=1}^{N} q_j \log(x_i - y_j), \qquad i = 1,\, 2,\, \dots,\, M$$

$$u_i = u(x_i), \qquad i = 1,\, 2,\, \dots,\, M$$

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

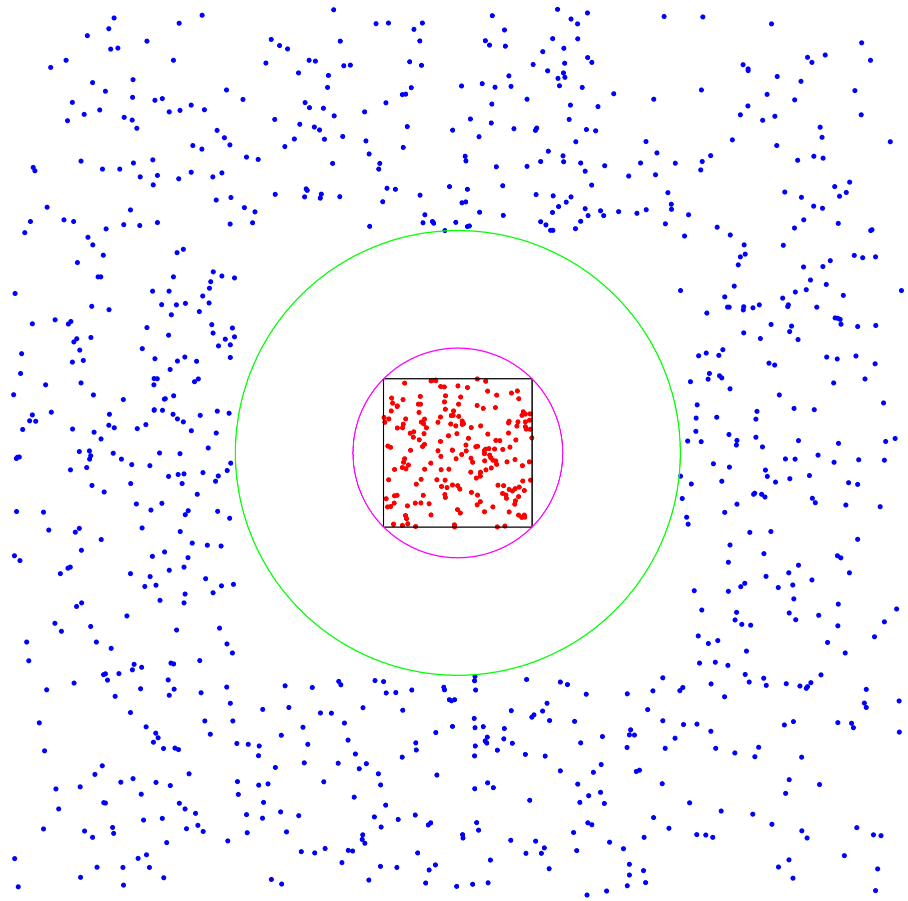

Multipole expansion:

Since u is harmonic outside the magenta circle, we know that

$$u(x) = \alpha_0 \log(x-c) + \sum_{p=1}^{\infty} \alpha_p \frac{1}{(x-c)^p}, \qquad \text{where } \alpha_p = -\frac{1}{p}\sum_{j=1}^{N} q_j\,(y_j - c)^p.$$

Special case: Sources and targets are separate

Charges q_j at locations $\{y_j\}_{j=1}^N$.

Potentials u_i at locations $\{x_i\}_{i=1}^M$.

$$u_i = \sum_{j=1}^{N} q_j \log(x_i - y_j), \qquad i = 1,\, 2,\, \dots,\, M$$

$$u_i = u(x_i), \qquad i = 1,\, 2,\, \dots,\, M$$

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

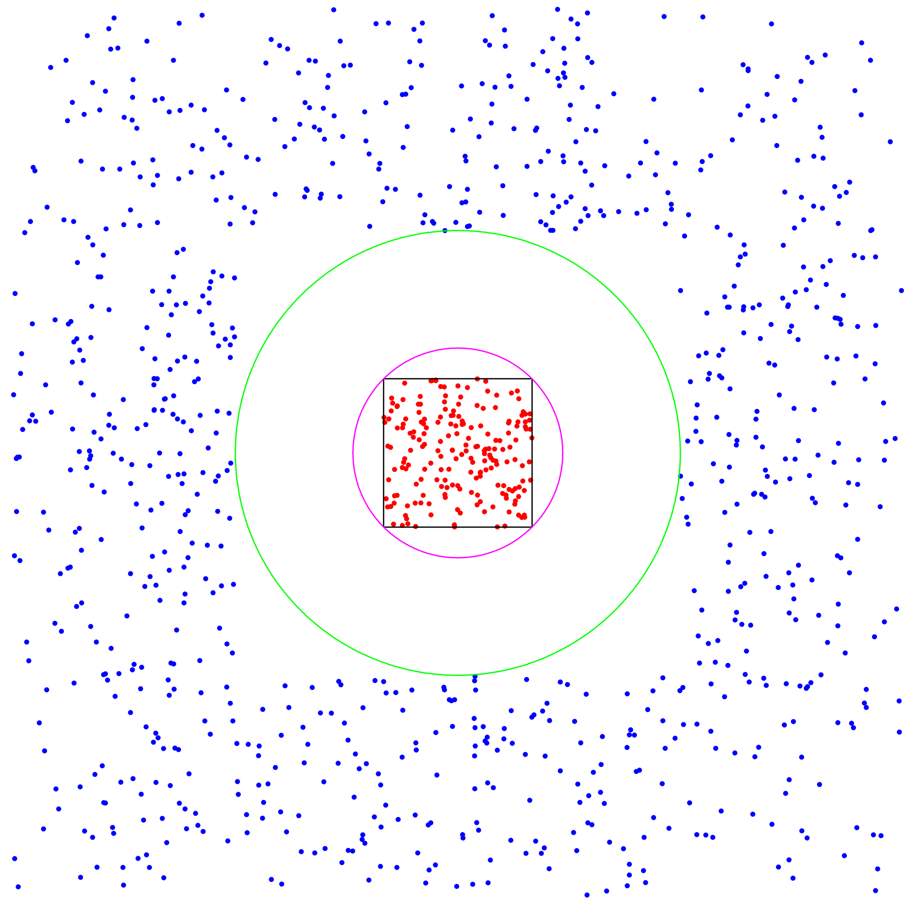

Multipole expansion — truncated to $P+1$ terms:

Since u is harmonic outside the magenta circle, we know that

$$u(x) = \alpha_0 \log(x - c) + \sum_{p=1}^{P} \alpha_p \frac{1}{(x-c)^p} + E_P, \qquad \text{where } \alpha_p = -\frac{1}{p} \sum_{j=1}^{N} q_j\, (y_j - c)^p.$$

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

Special case: Sources and targets are separate

Charges q_j at locations $\{y_j\}_{j=1}^{N}$.

Potentials u_i at locations $\{x_i\}_{i=1}^{M}$.

$$u_i = \sum_{j=1}^{N} q_j \log(x_i - y_j), \qquad i = 1,\, 2,\, \dots,\, M$$

$$u_i = u(x_i), \qquad i = 1,\, 2,\, \dots,\, M$$

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

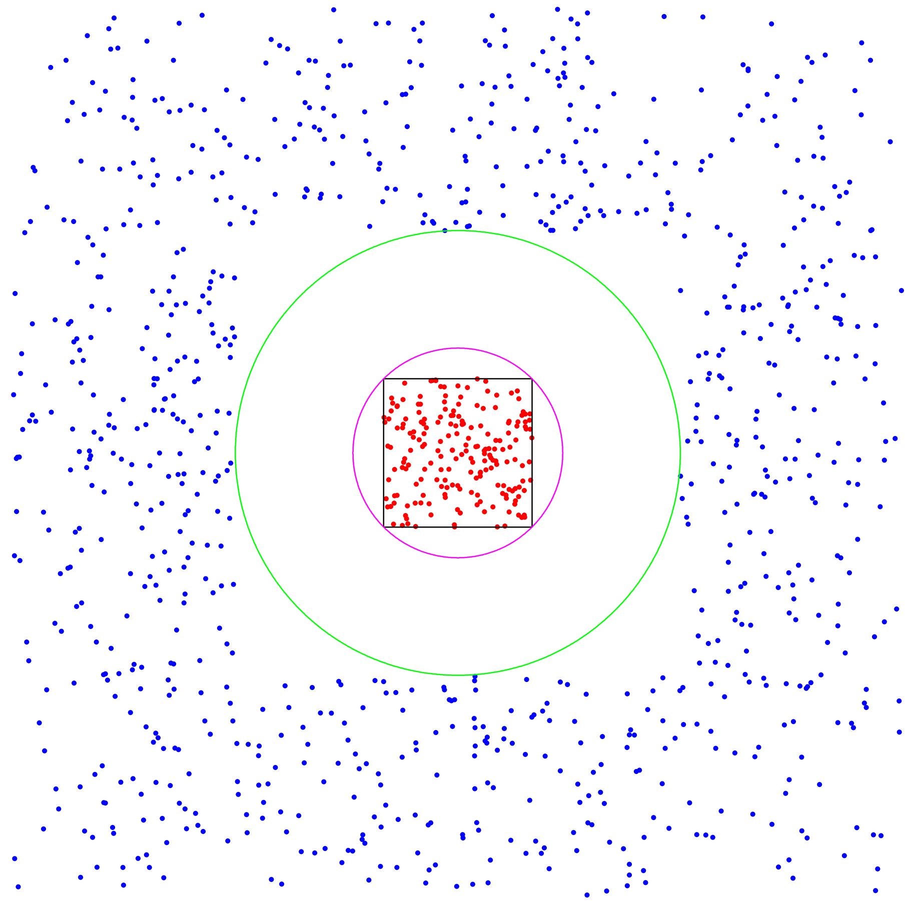

Multipole expansion — truncated to $P+1$ terms — costs:

Evaluate $\alpha_p = -\frac{1}{p}\sum_{j=1}^{N} q_j\,(y_j - c)^p$ for $p = 0,\, 1,\, \dots,\, P$ — *cost is $O(N\,P)$.*

Evaluate $u_i = \alpha_0 \log(x - c) + \sum_{p=1}^{P} \alpha_p \frac{1}{(x-c)^p}$ — *cost is $O(M\,P)$.*

We write the map in matrix form as $||\mathsf{A} - \tilde{\mathsf{B}}\,\mathsf{C}|| \leq \varepsilon$, or

$$\begin{tikzcd} \boldsymbol{q} \arrow[r, "\mathsf{A}"] \arrow[d, "\mathsf{C}"'] & \boldsymbol{u} \\ \hat{\boldsymbol{q}} \arrow[ur, "\tilde{\mathsf{B}}"'] & \end{tikzcd}$$

where

$\boldsymbol{q} = [q_j]_{j=1}^{N} \in \mathbb{R}^N$ is the vector of *sources*

$\boldsymbol{u} = [u_i]_{i=1}^{M} \in \mathbb{C}^M$ is the vector of *potentials*

$\hat{\boldsymbol{q}} = [\hat{q}_p]_{p=0}^{N} \in \mathbb{C}^{P+1}$ is the *outgoing expansion* (the "multipole coefficients")

$\mathsf{A} \in \mathbb{C}^{M\times N}$

$\mathsf{C} \in \mathbb{C}^{(P+1)\times N}$ is the *outgoing-from-sources* map

$\tilde{\mathsf{B}} \in \mathbb{C}^{M\times(P+1)}$ is the *targets-from-outgoing* map.

$$\mathsf{A}_{ij} = \log(x_i - y_j)$$

$$\mathsf{C}_{pj} = \begin{cases} 1 & p = 0 \\ -(1/p)\,(y_j - c)^p & p \neq 0 \end{cases}$$

$$\tilde{\mathsf{B}}_{ip} = \begin{cases} \log(x_i - c) & p = 0 \\ 1/(x_i - c)^p & p \neq 0 \end{cases}$$

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

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

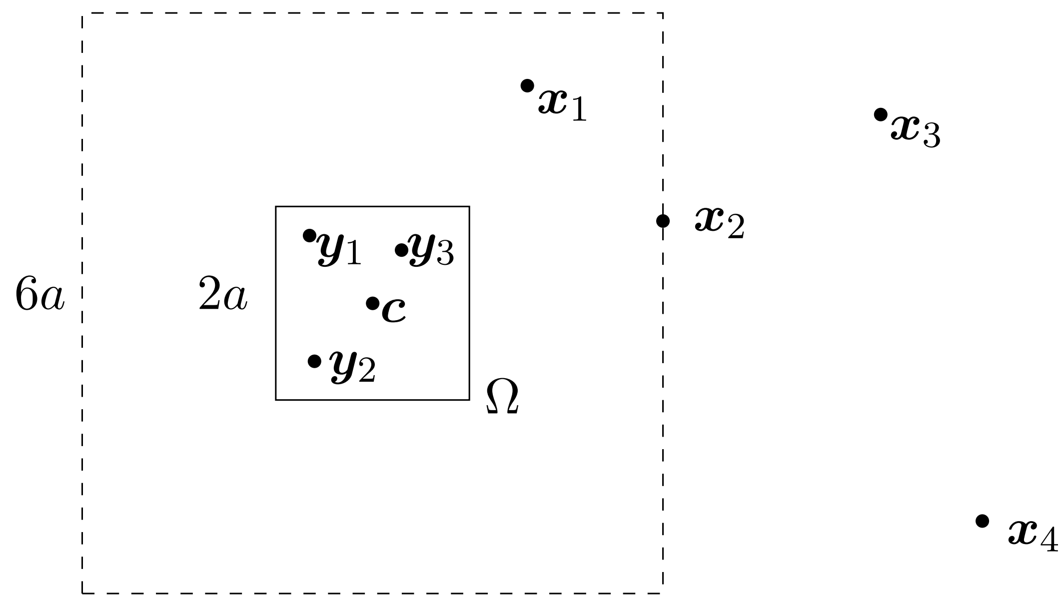

Any point on or outside of the dashed square is well-separated from Ω*.*
Consequently, the points $\boldsymbol{x}_2$*,* $\boldsymbol{x}_3$*, and* $\boldsymbol{x}_4$ *are well-separated from* Ω*, but* $\boldsymbol{x}_1$ *is not.*

Single-level Barnes-Hut

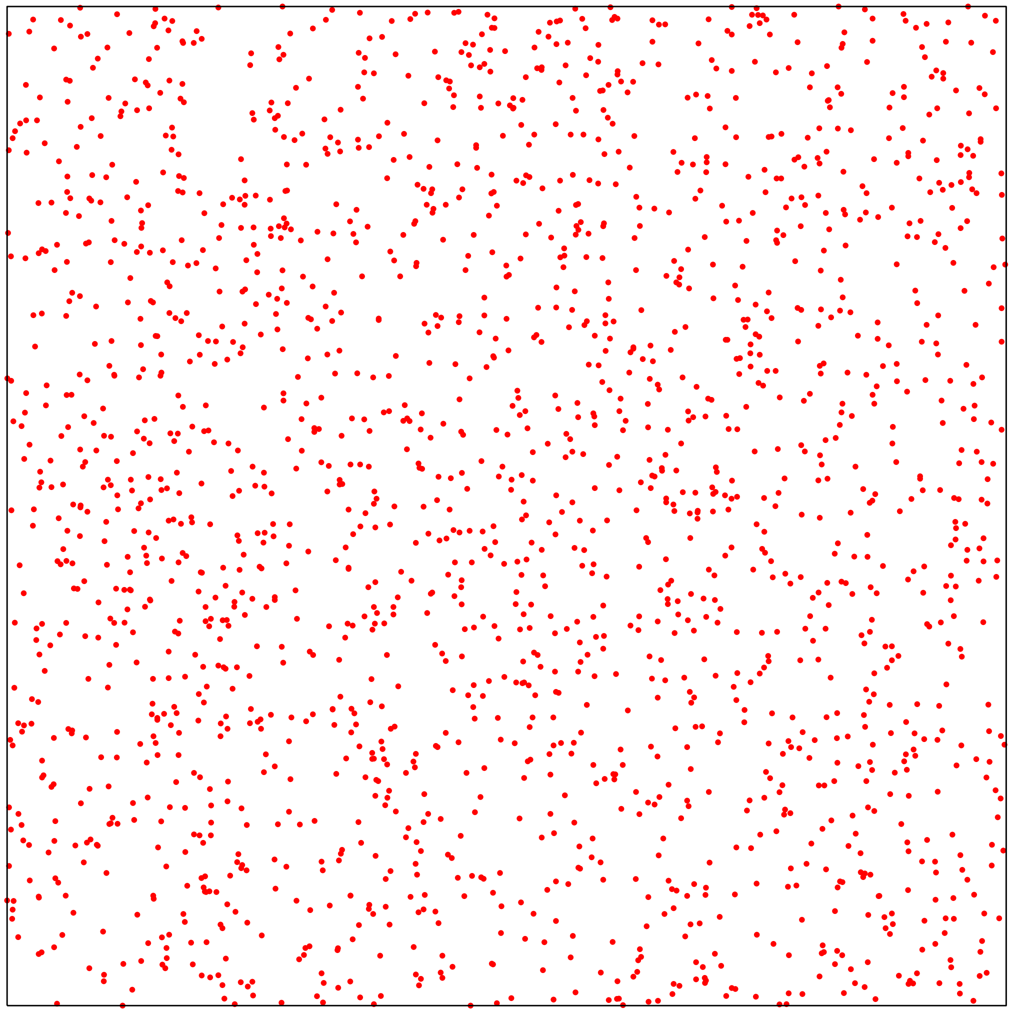

Suppose that we seek to evaluate all pairwise interactions between the N particles shown above.

We suppose the particles are more or less evenly distributed for now.

Single-level Barnes-Hut

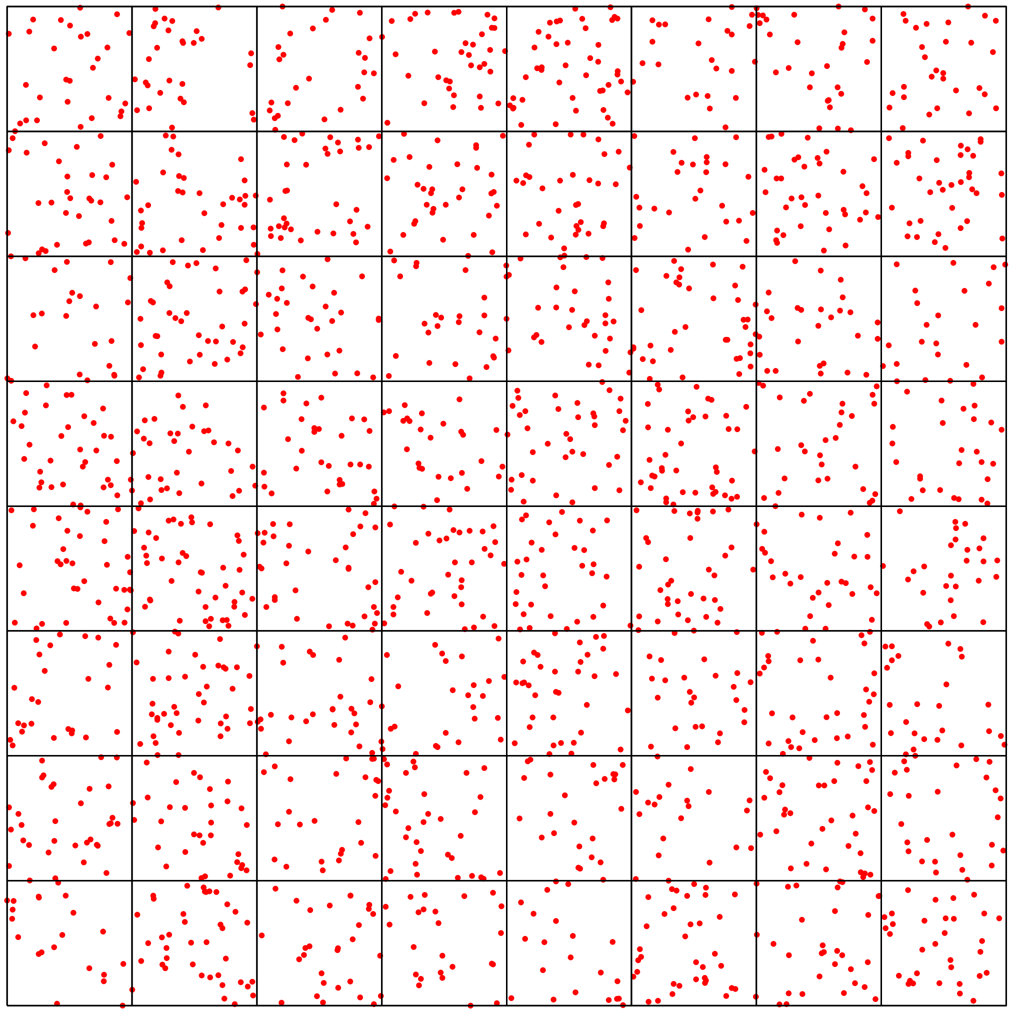

Place a square grid of boxes on top of the computational box.

Assume each box holds about N_{box} particles (so there are about N/N_{box} boxes).

Single-level Barnes-Hut

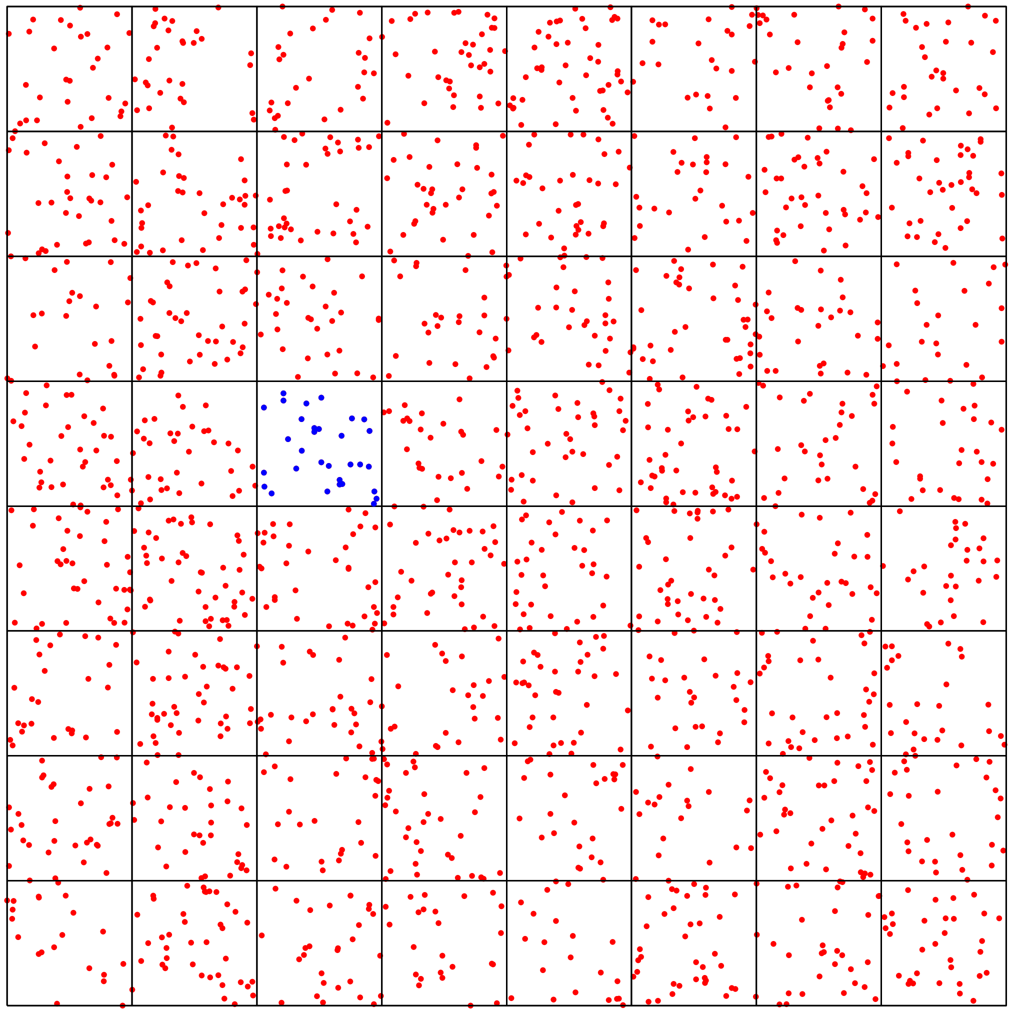

How do you evaluate the potentials at the blue locations?

Single-level Barnes-Hut

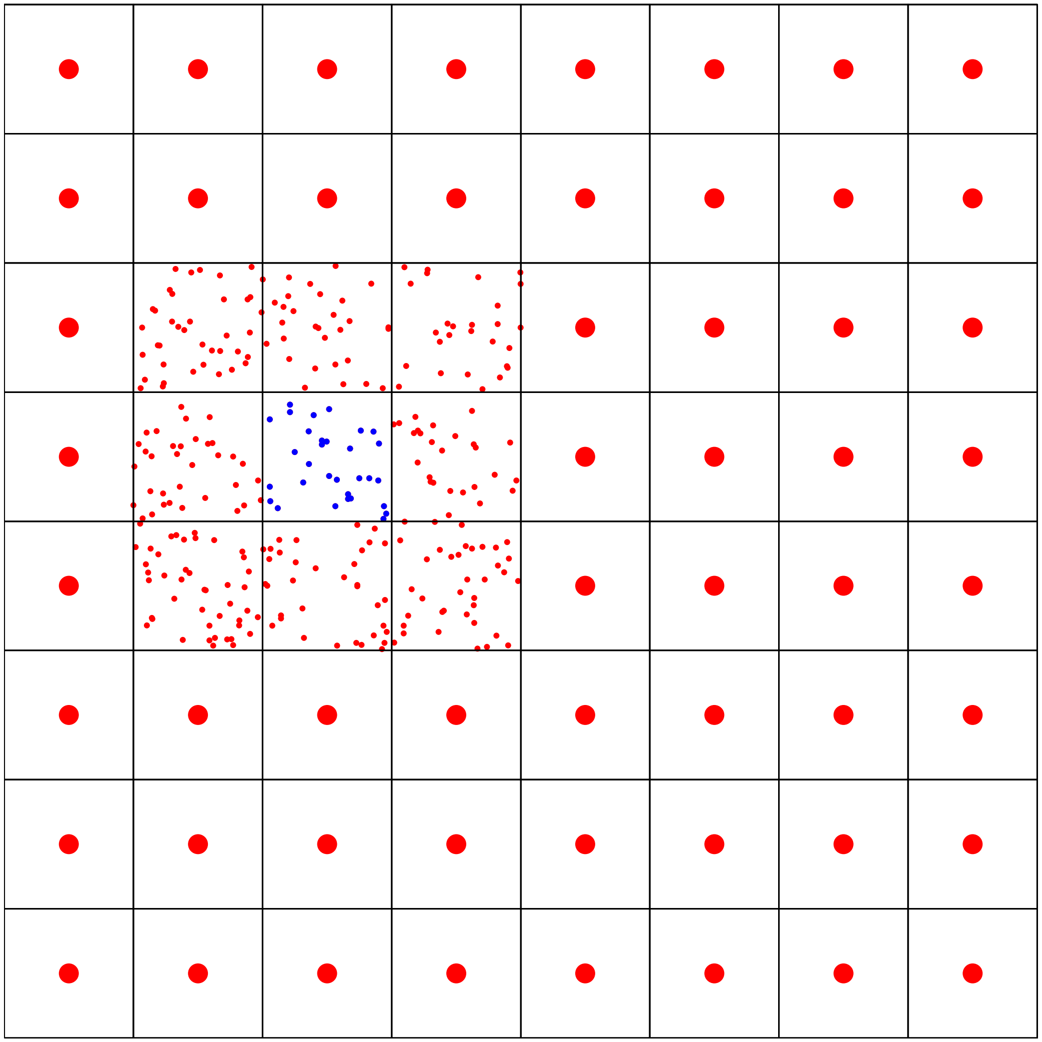

How do you evaluate the potentials at the blue locations?

Directly evaluate interactions with the particles close by.

For all long-distance interactions, use the out-going expansion!

Cost of the single-level Barnes-Hut algorithm:

Let $N_{\rm box}$ 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 out-going expansions:	$((N/N_{\rm box}) - 9)\,P$
Step 3:	Evaluate potentials from close particles:	$9\,N_{\rm box}$

Using that P is a smallish constant we find

$$\text{cost} \sim \frac{N^2}{N_{\rm box}} + N\,N_{\rm box}.$$

Set $N_{\rm box} \sim N^{1/2}$ 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	$\boxed{21}$	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 $\mathcal{L}_\tau^{(\mathrm{nei})}$ denote its neighbors, and $\mathcal{L}_\tau^{(\mathrm{far})}$ 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:

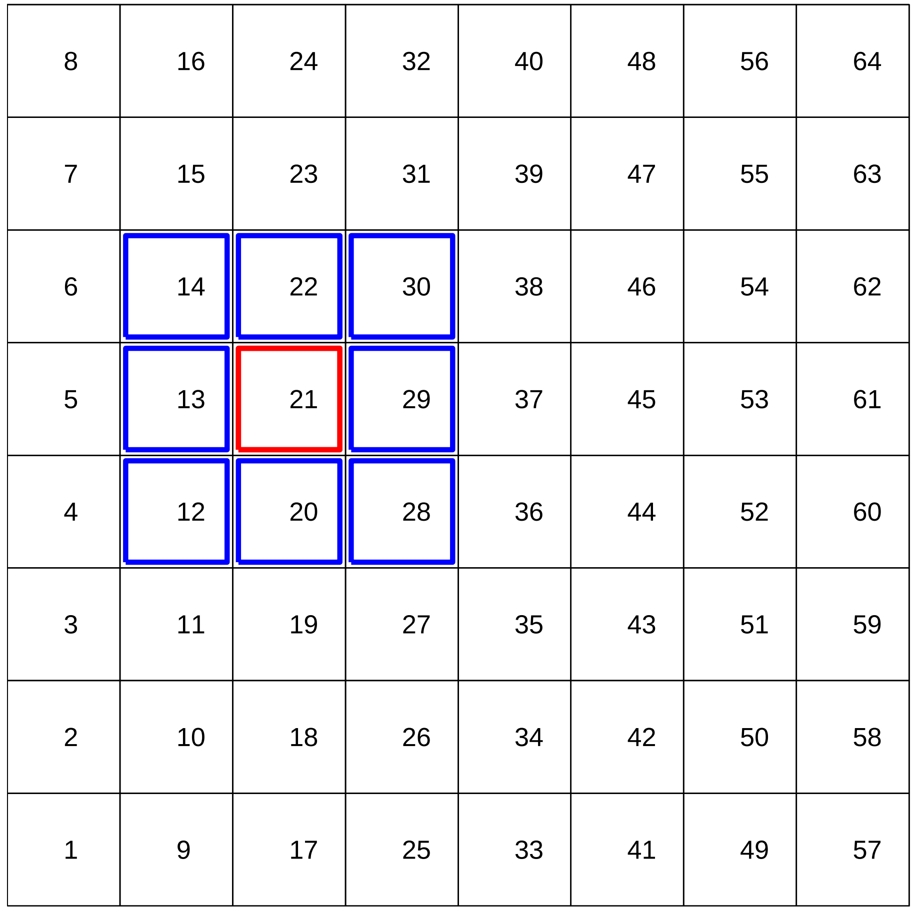

For a box τ, let $\mathcal{L}_\tau^{(\mathrm{nei})}$ denote its neighbors, and $\mathcal{L}_\tau^{(\mathrm{far})}$ denote the remaining boxes.

The box $\tau = 21$ is marked with red.
$\mathcal{L}_{21}^{(\mathrm{nei})} = \{12, 13, 14, 20, 22, 28, 29, 30\}$ are the blue boxes.

Notation — single level Barnes-Hut:

Partition the box Ω into smaller boxes and label them:

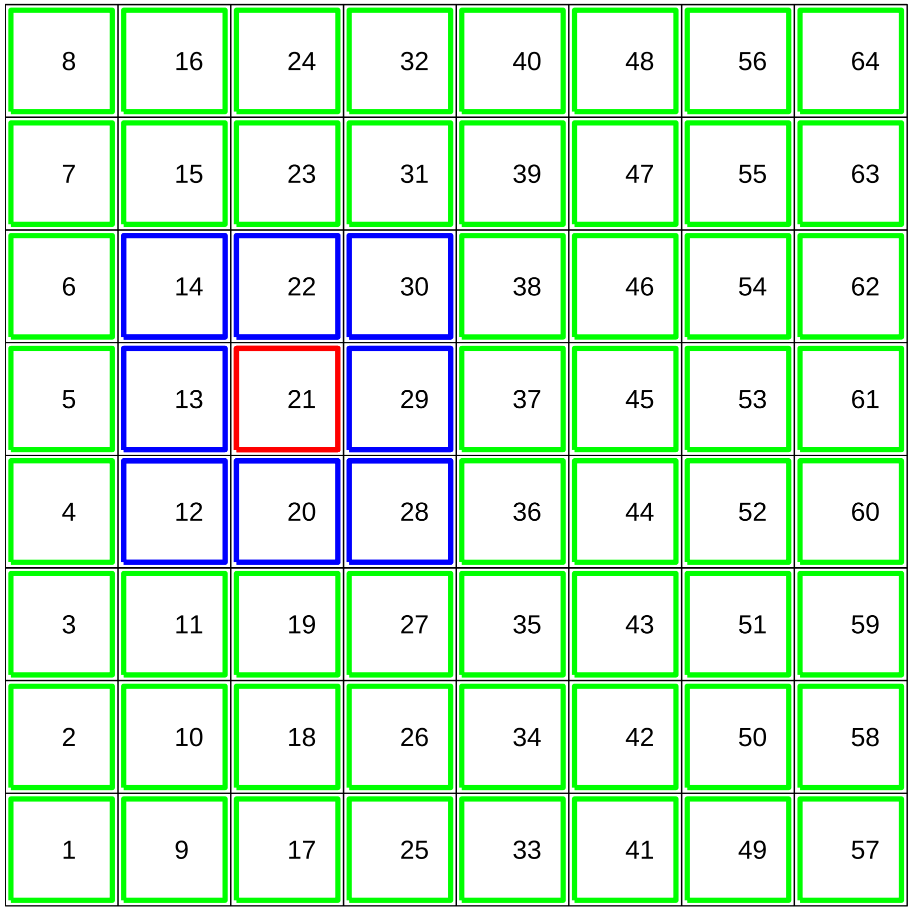

For a box τ, let $\mathcal{L}_\tau^{(\text{nei})}$ denote its neighbors, and $\mathcal{L}_\tau^{(\text{far})}$ denote the remaining boxes.

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

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

$$j \in J_\tau \qquad \Leftrightarrow \qquad \boldsymbol{x}_j \text{ is in box } \tau.$$

Set

$\boldsymbol{q}_\tau = \boldsymbol{q}(J_\tau)$ — the sources in τ

$\hat{\boldsymbol{q}}_\tau \in \mathbb{C}^{P+1}$ — the outgoing expansion for τ (the "multipole coefficients")

$\boldsymbol{u}_\tau = \boldsymbol{u}(J_\tau)$ — the potentials in τ

Then

$$\begin{tikzcd} \boldsymbol{q}_\tau \arrow[r, "\mathsf{A}_{\sigma,\tau}"] \arrow[d, "\mathsf{C}_\tau"'] & \boldsymbol{u}_\sigma \\ \hat{\boldsymbol{q}}_\tau \arrow[ur, "\tilde{\mathsf{B}}_{\sigma,\tau}"'] & \end{tikzcd}$$

Recall: P is an integer parameter balancing cost versus accuracy:

$$P \text{ large} \qquad \Rightarrow \qquad \text{high accuracy and high cost}$$

Single-level Barnes-Hut

Compute the outgoing expansions on all boxes:

loop over all boxes τ

$\quad \hat{\boldsymbol{q}}_\tau = \mathsf{C}_\tau\, \boldsymbol{q}(J_\tau)$

end loop

Evaluate the far field potentials.
Each box τ aggregates the contributions from all well-separated boxes:

$\boldsymbol{u} = 0$

loop over all boxes τ

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

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

$\quad$ **end loop**

end loop

Evaluate the near field interactions:

loop over all leaf boxes τ

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

end

To get the asymptotic cost down further, we need a hierarchical tree structure on the computational domain:

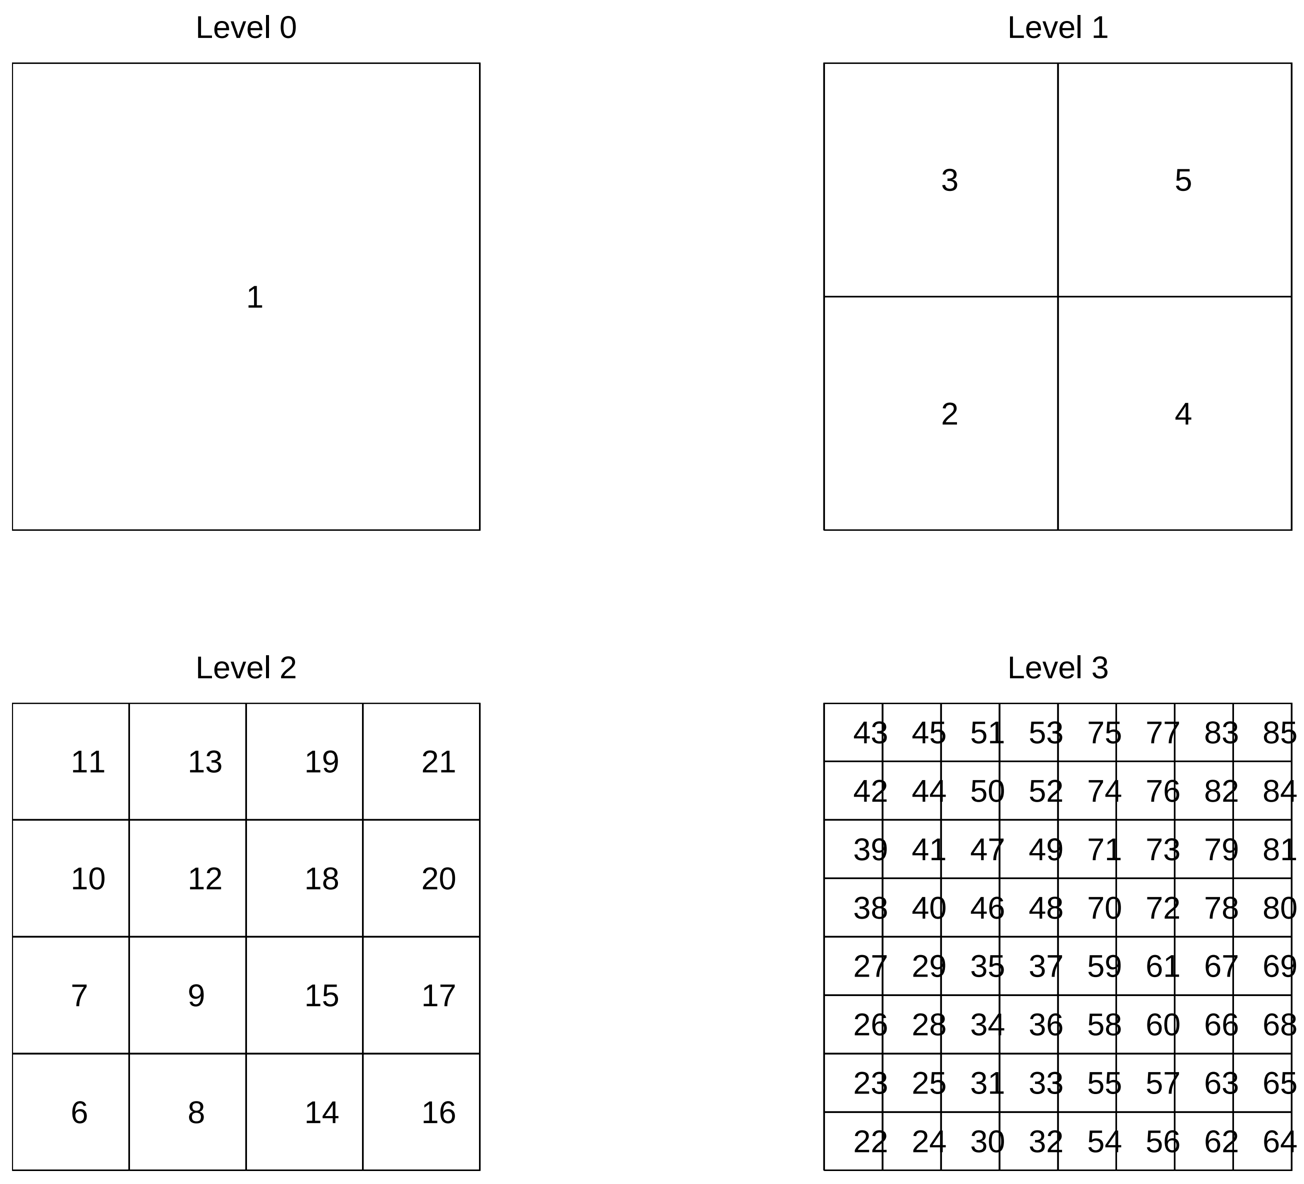

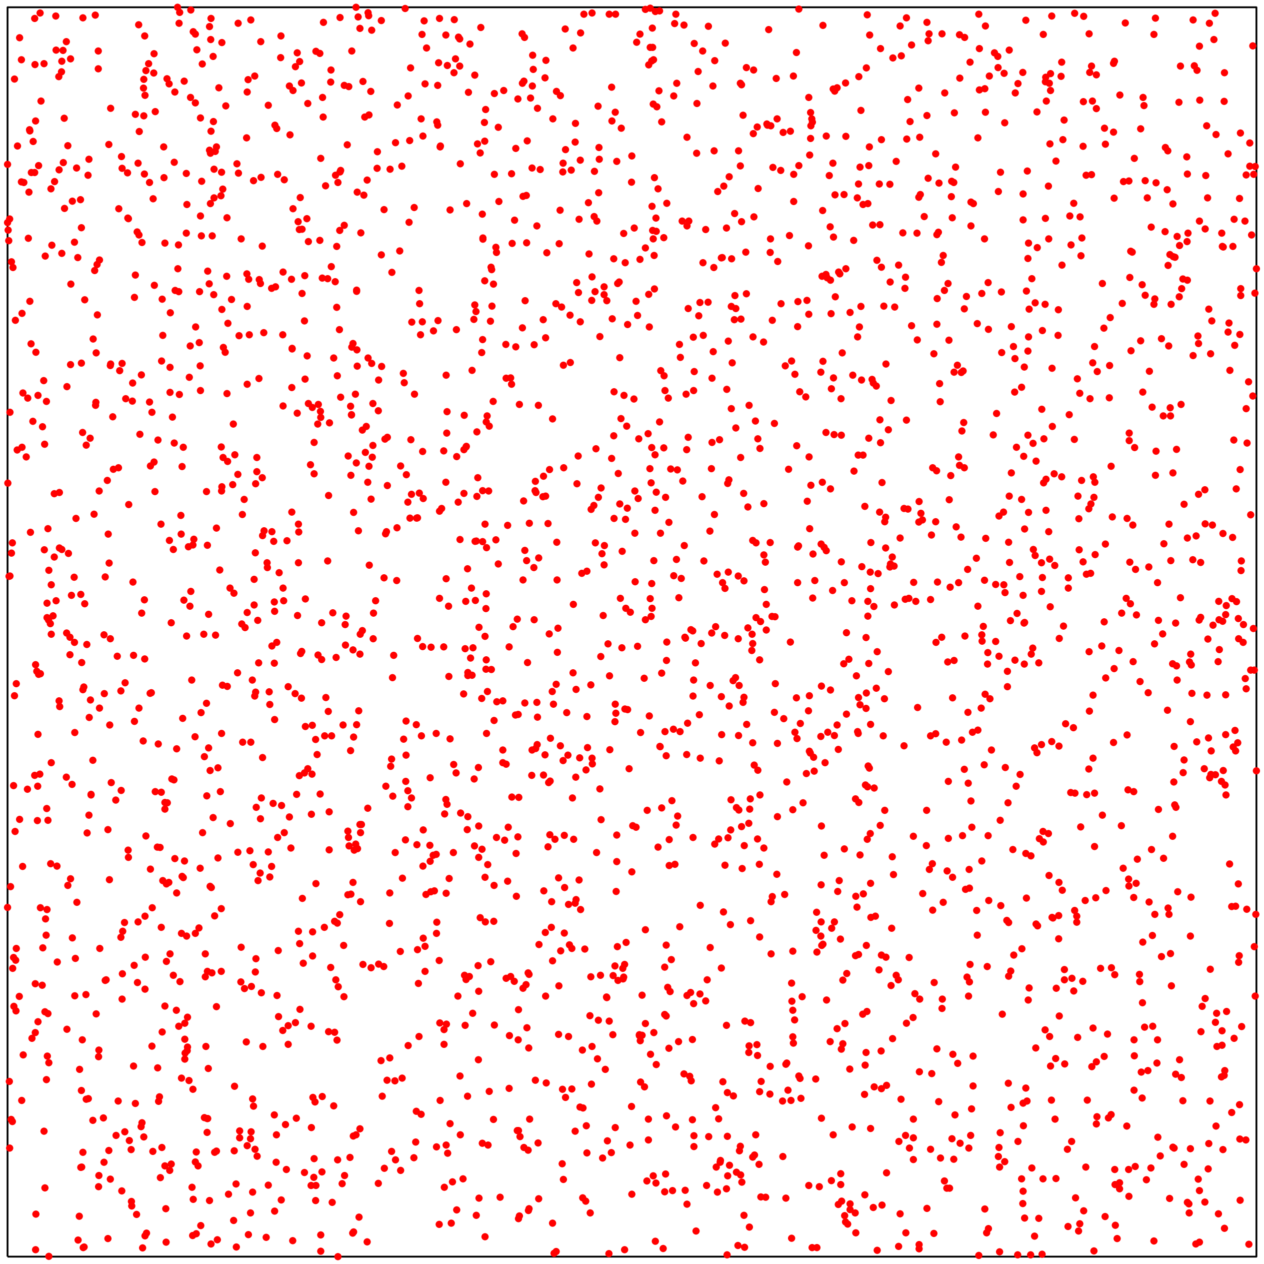

We are given N locations and seek the potential at each location.

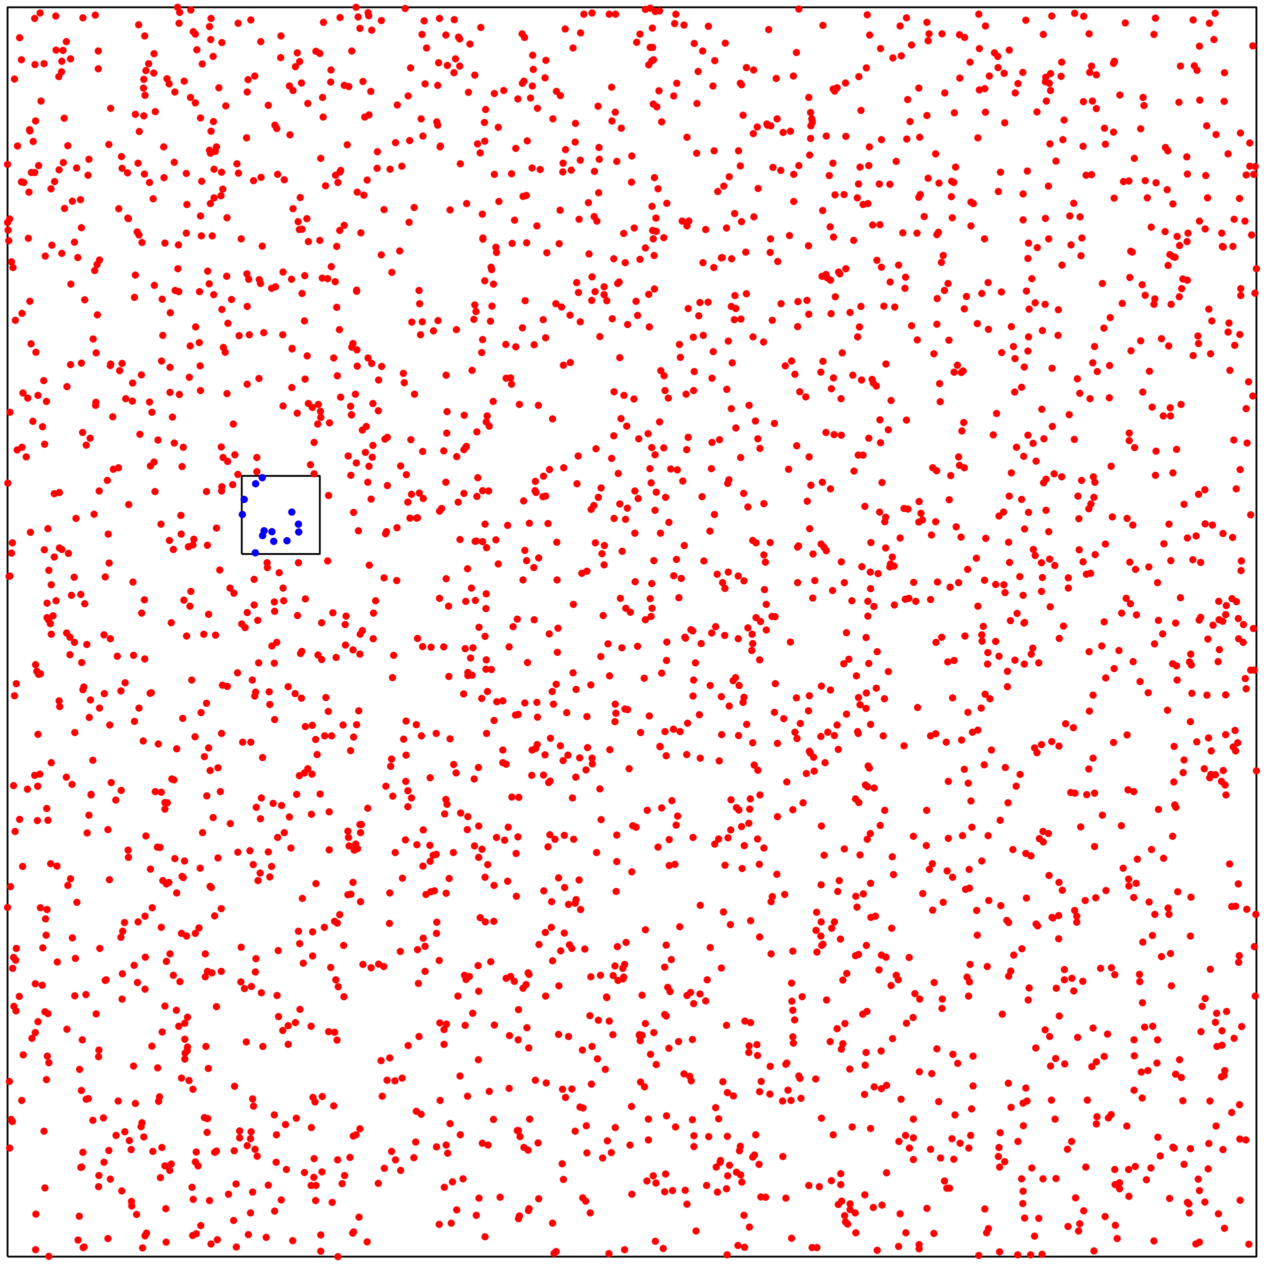

How do you find the potential at the locations marked in blue?

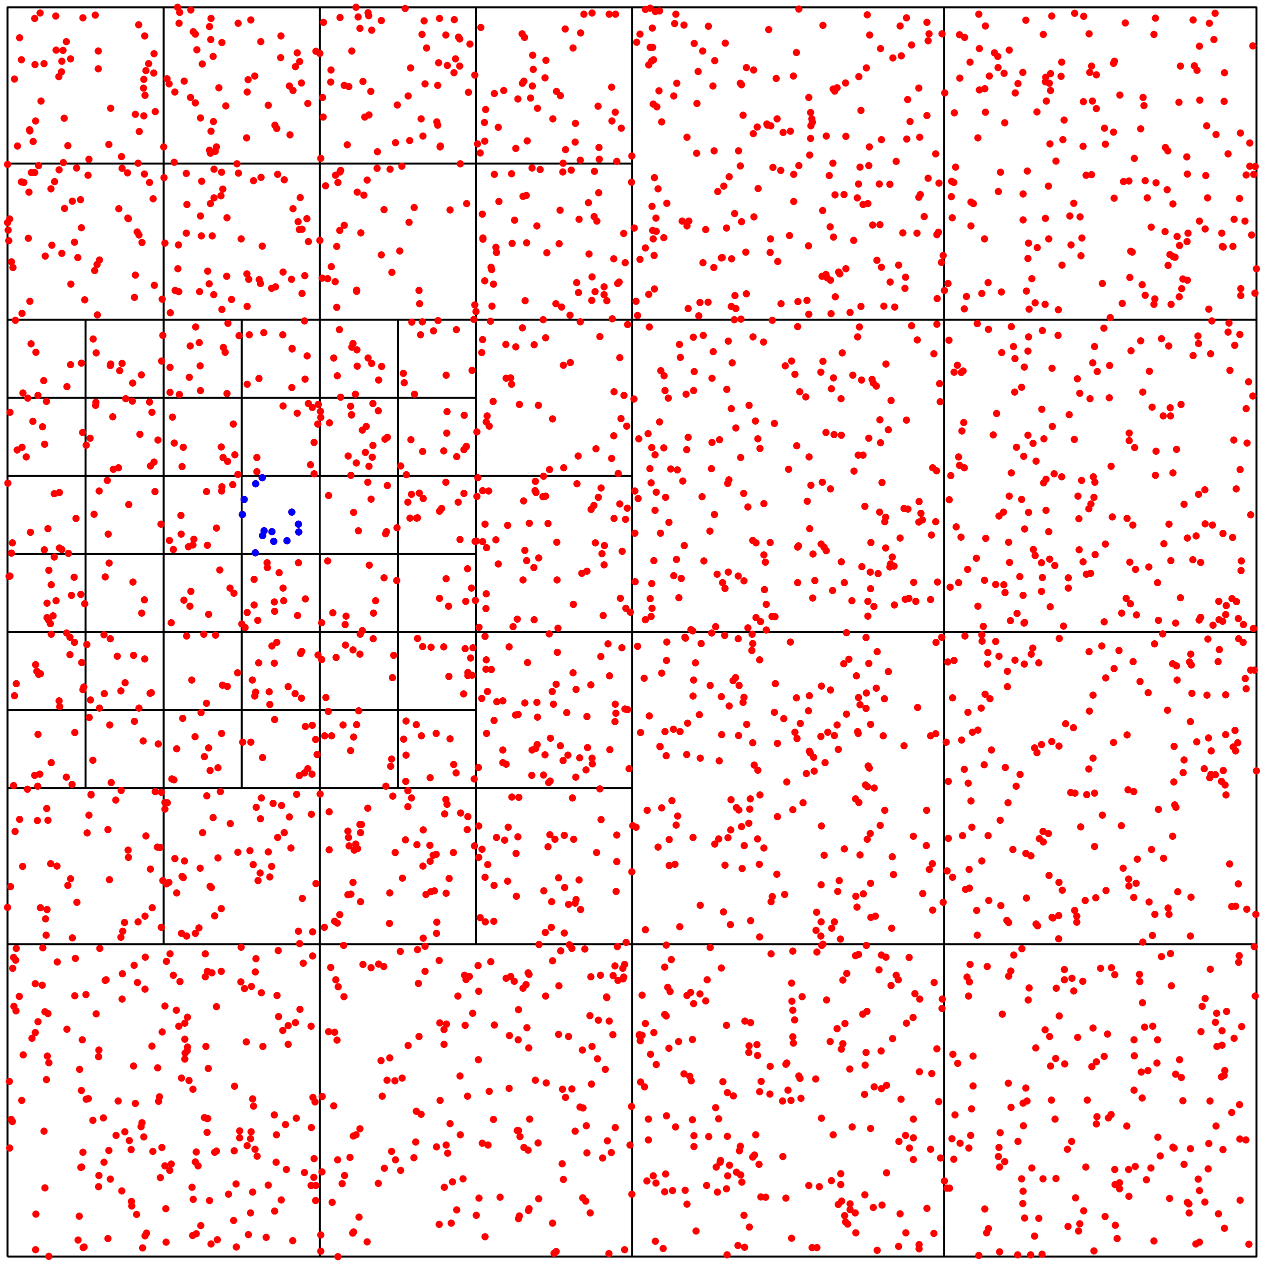

Tessellate the remainder of the domain using as large boxes as you can, with the constraint that *the target box has to be well-separated from every box that is used.*

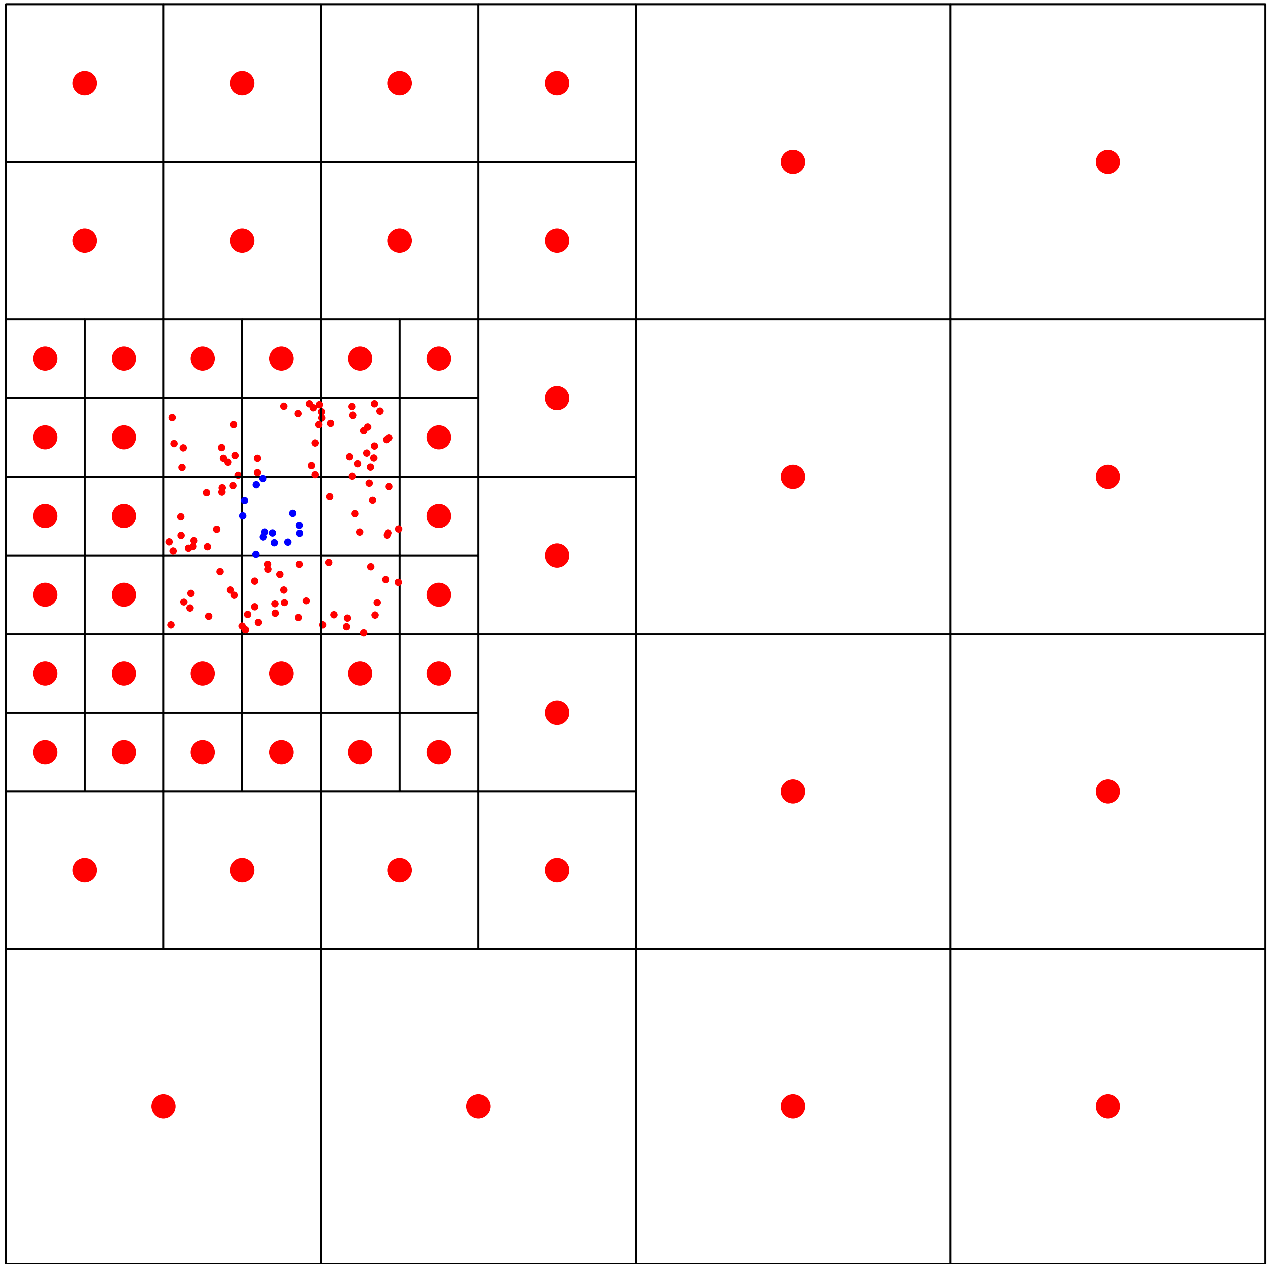

Then replace the original charges in each well-separated box by the corresponding multipole expansion.

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 $N_{\rm box}$ particles in each box so that $L \approx \log_4(N/N_{\rm box})$.

We let $N_{\rm box}$ is a fixed number (say $N_{\rm box} \approx 200$) so $L \sim \log(N)$.

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

For each particle $\boldsymbol{x}_i$, we need to do the following:

Step 1: Evaluate the P outgoing moments for the L boxes holding x_i:	$L\,P$
Step 2: Evaluate potentials from out-going expansions:	$27\,L\,P$
Step 3: Evaluate potentials from neighbors:	$9\,N_{\rm box}$

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

$\qquad \hat{\boldsymbol{q}}_\tau = \mathsf{C}_\tau\, \boldsymbol{q}(J_\tau)$

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:

$\boldsymbol{u} = 0$

loop over all boxes τ

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

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

$\qquad$ **end loop**

end loop

Evaluate the near field interactions:

loop over all leaf boxes τ

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

end

In the Barnes-Hut method, there is a need to construct for each box τ, the list

$$\mathcal{L}_\tau^{(\mathrm{BH})}$$

which identifies a "minimal" tessellation of Ω into boxes from which τ is well-separated.

Example: The box τ is marked with blue.

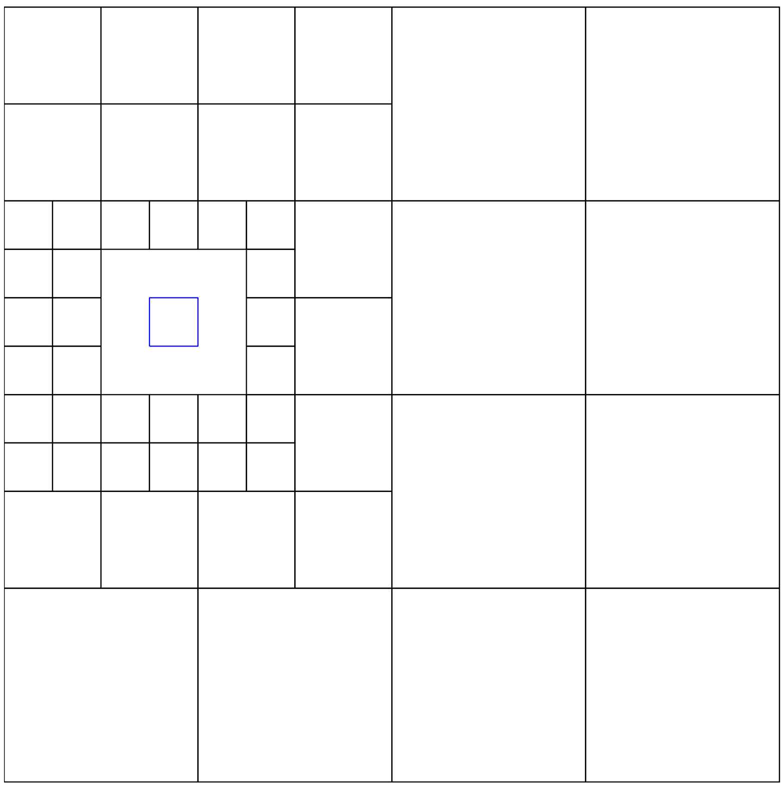

$$\mathcal{L}_\tau^{(\mathrm{BH})} = \{???\}$$

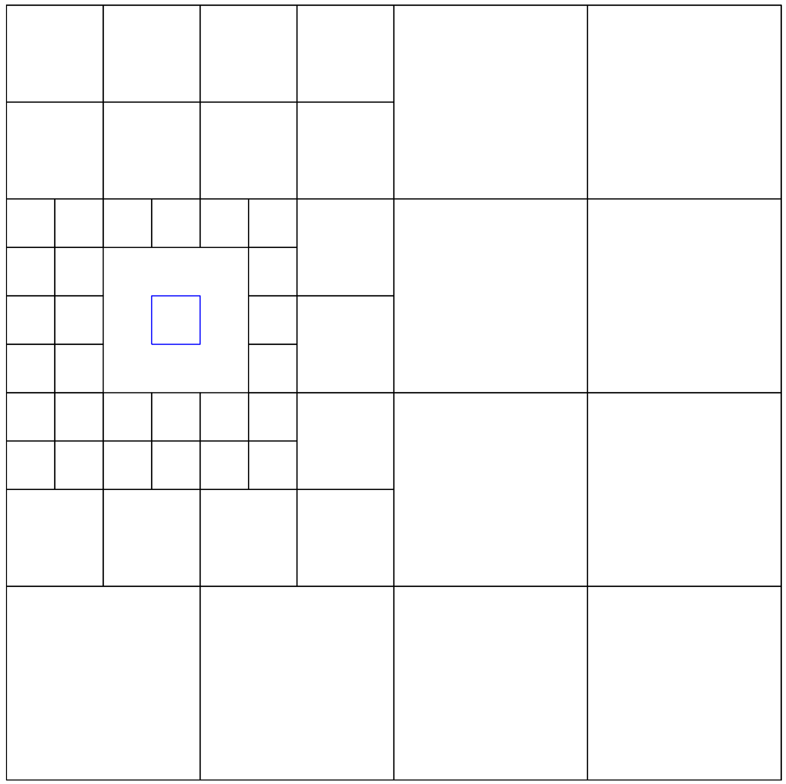

How do you find the tessellation?

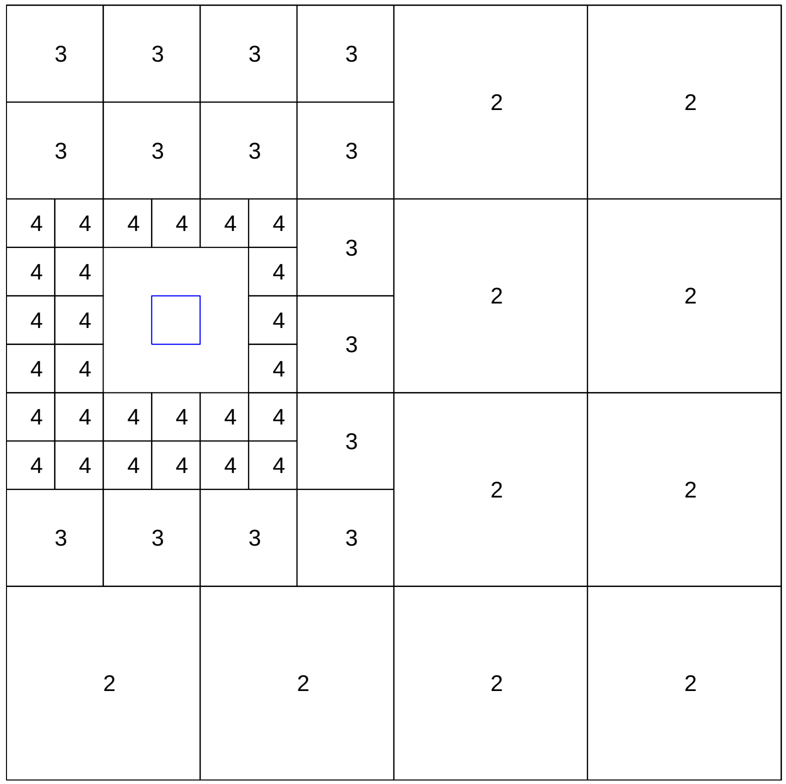

How do you find the tessellation?

Hint: Consider which "level" each box belongs to.

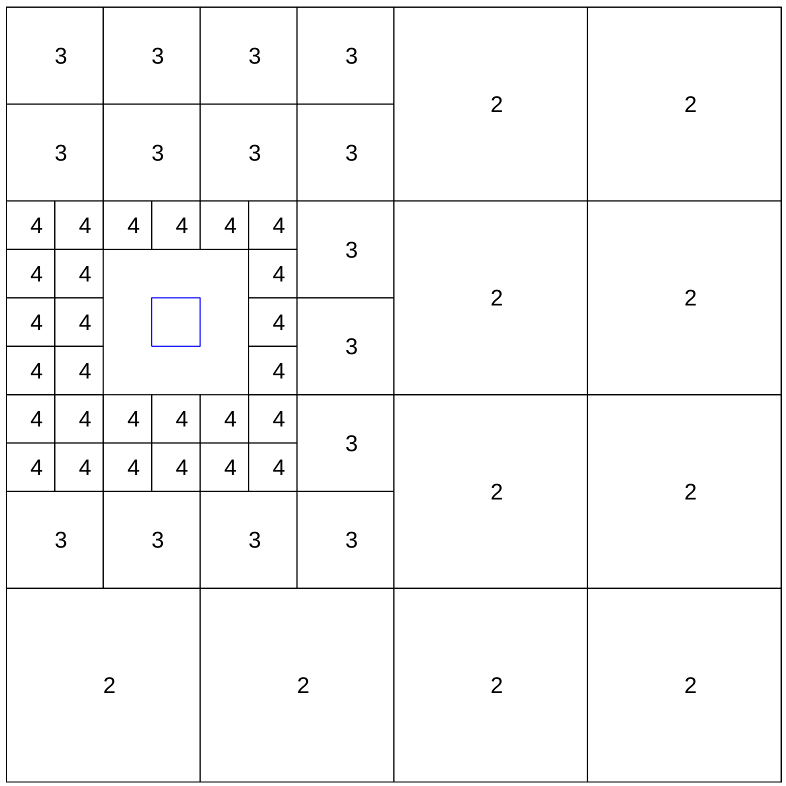

Level 4

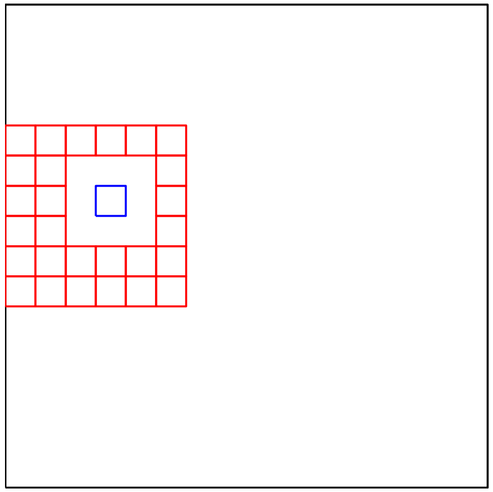

Level 3

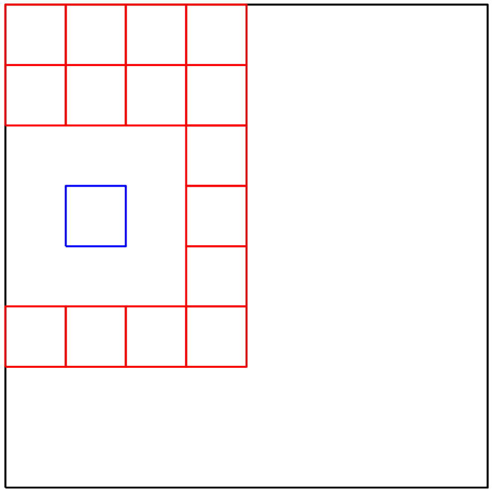

Level 2

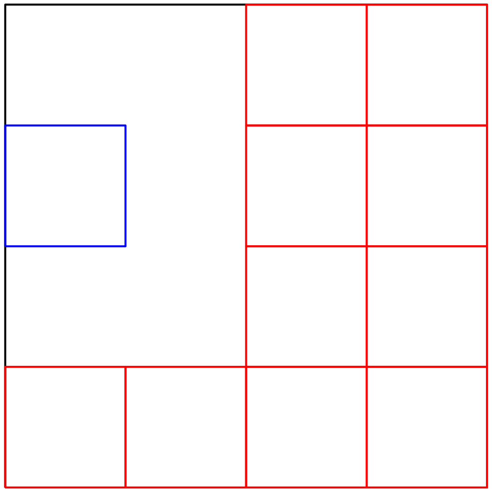

Definition:

For a box τ, define its *interaction list* $\mathcal{L}_\tau^{(\mathrm{int})}$ 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.

Example: Consider box $\tau = 10$.

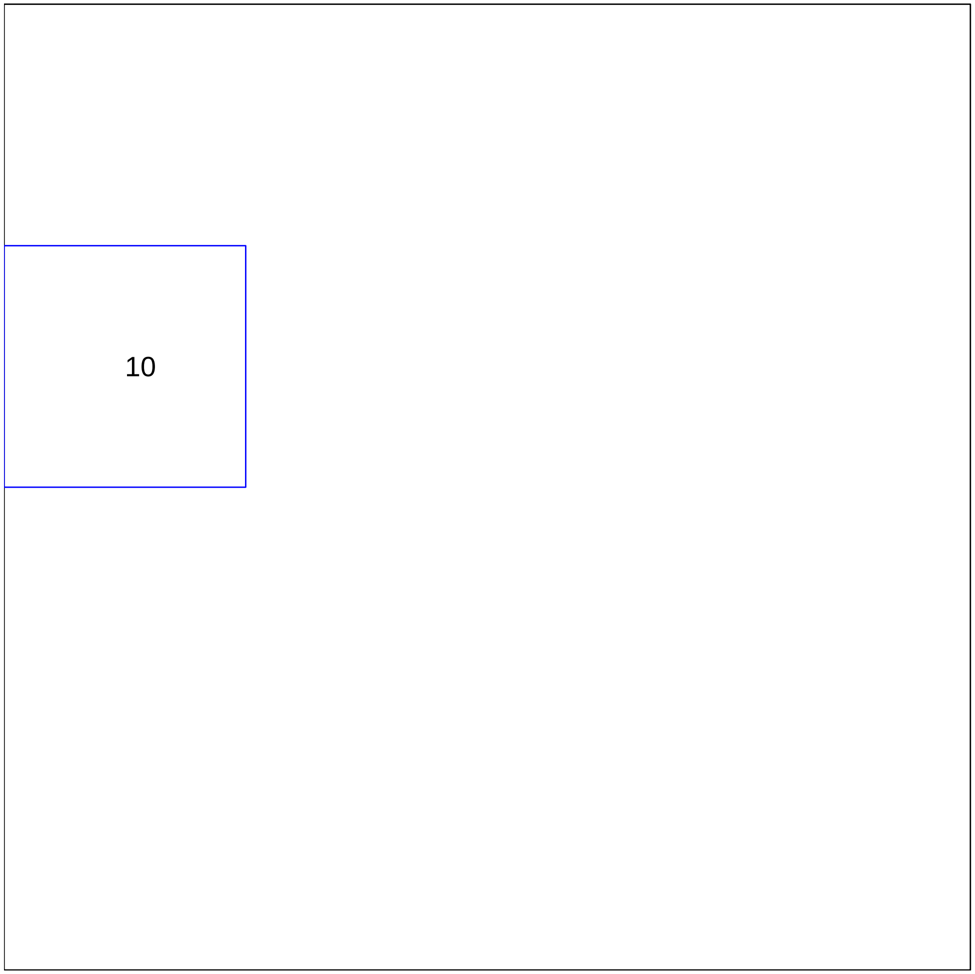

Example: Consider box $\tau = 10$.

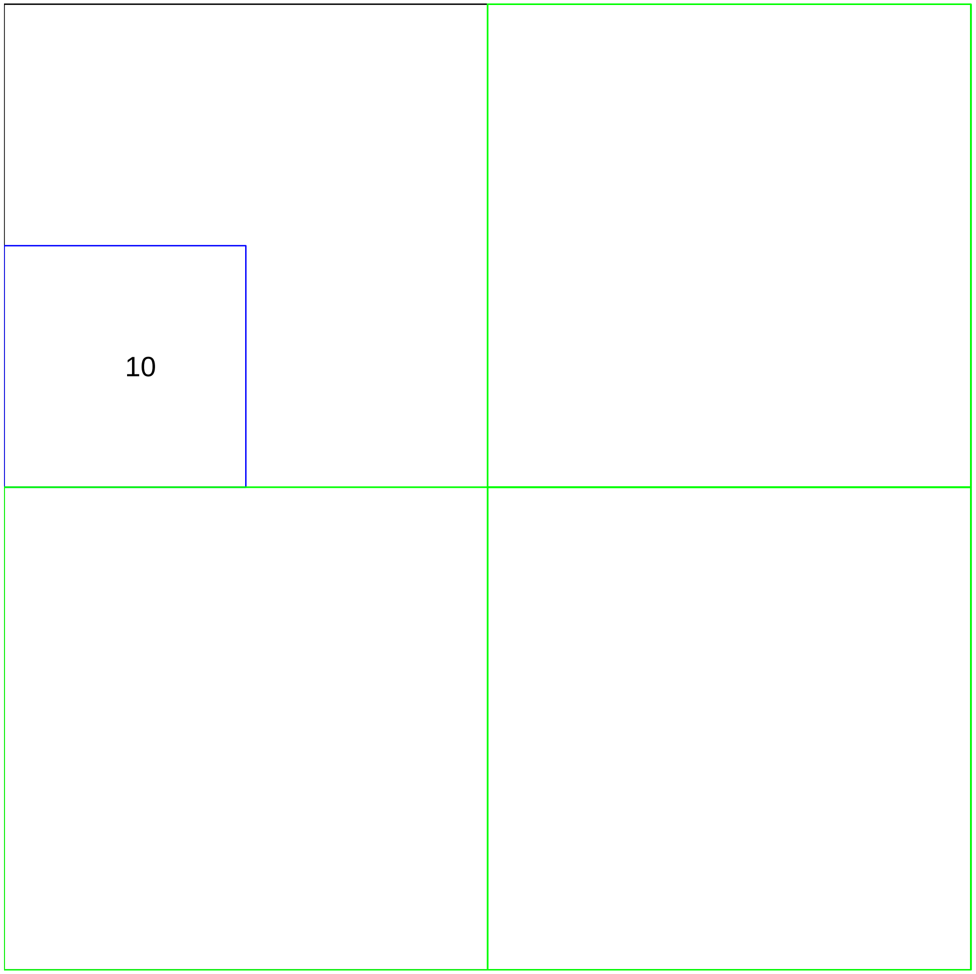

The green boxes are the neighbors of τ's parents.

Example: Consider box $\tau = 10$.

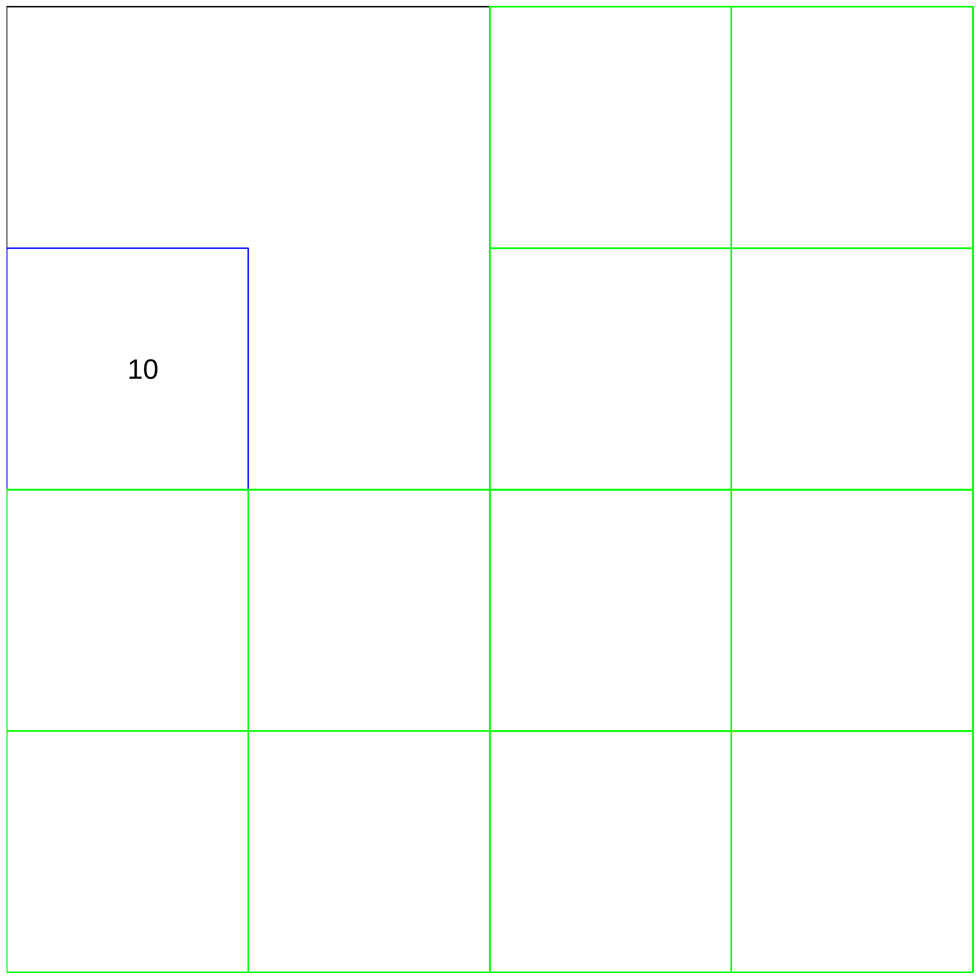

The green boxes are the children of the neighbors of τ's parents.

Example: Consider box $\tau = 10$.

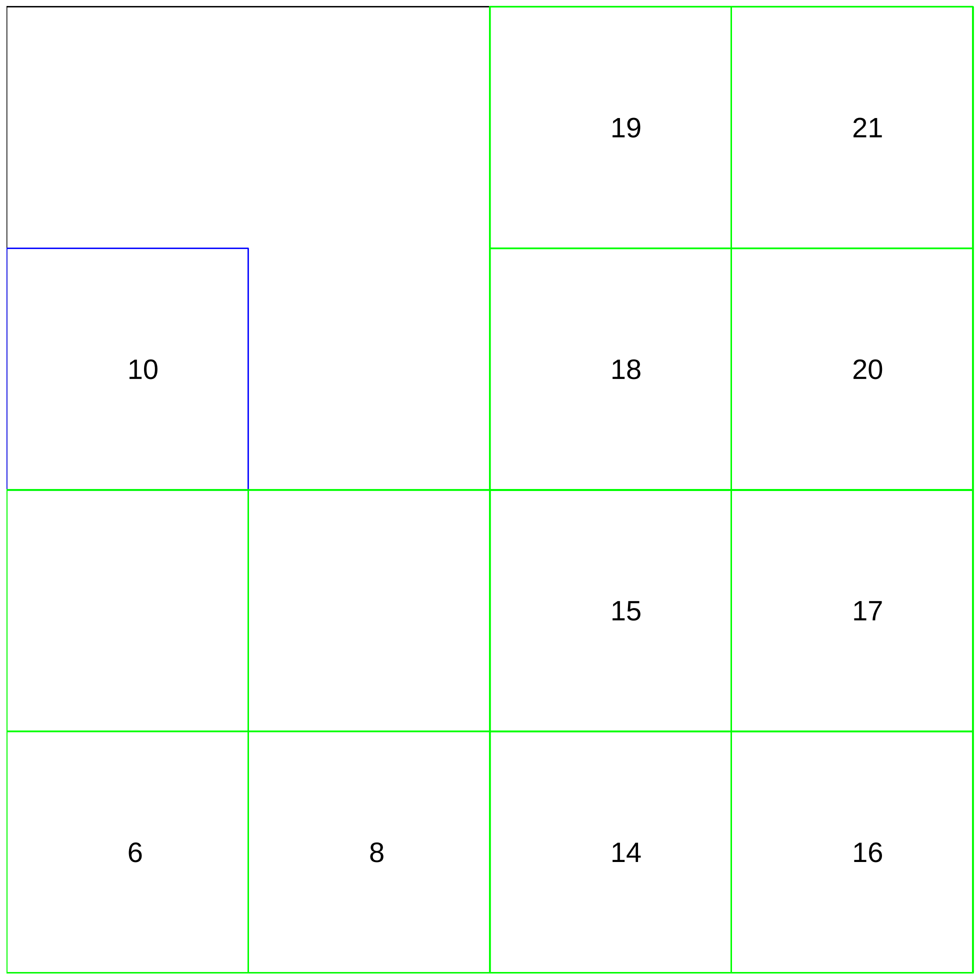

The green boxes are the children of the neighbors of τ's parents.

Only some of them are well-separated from τ.

We find: $\mathcal{L}_\tau^{(\mathrm{int})} = \{6, 8, 14, 15, 16, 17, 18, 19, 20, 21\}$

Example: Consider box $\tau = 40$.

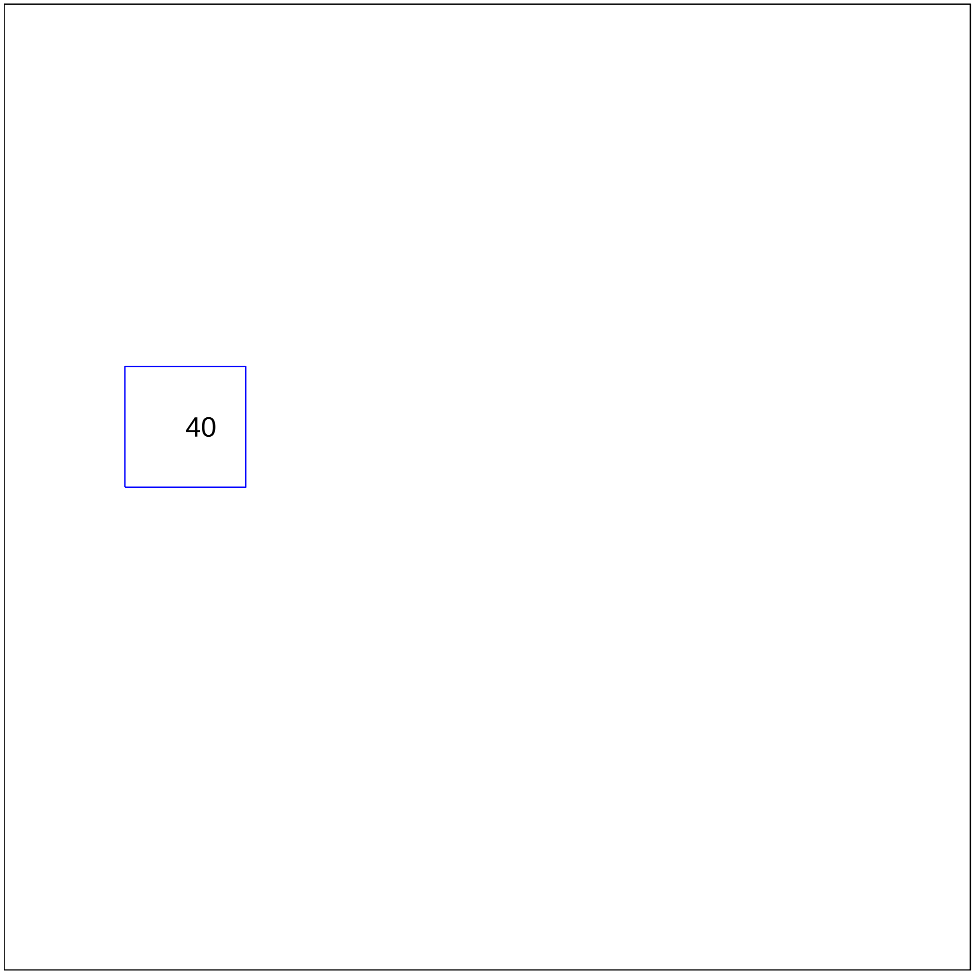

Example: Consider box $\tau = 40$.

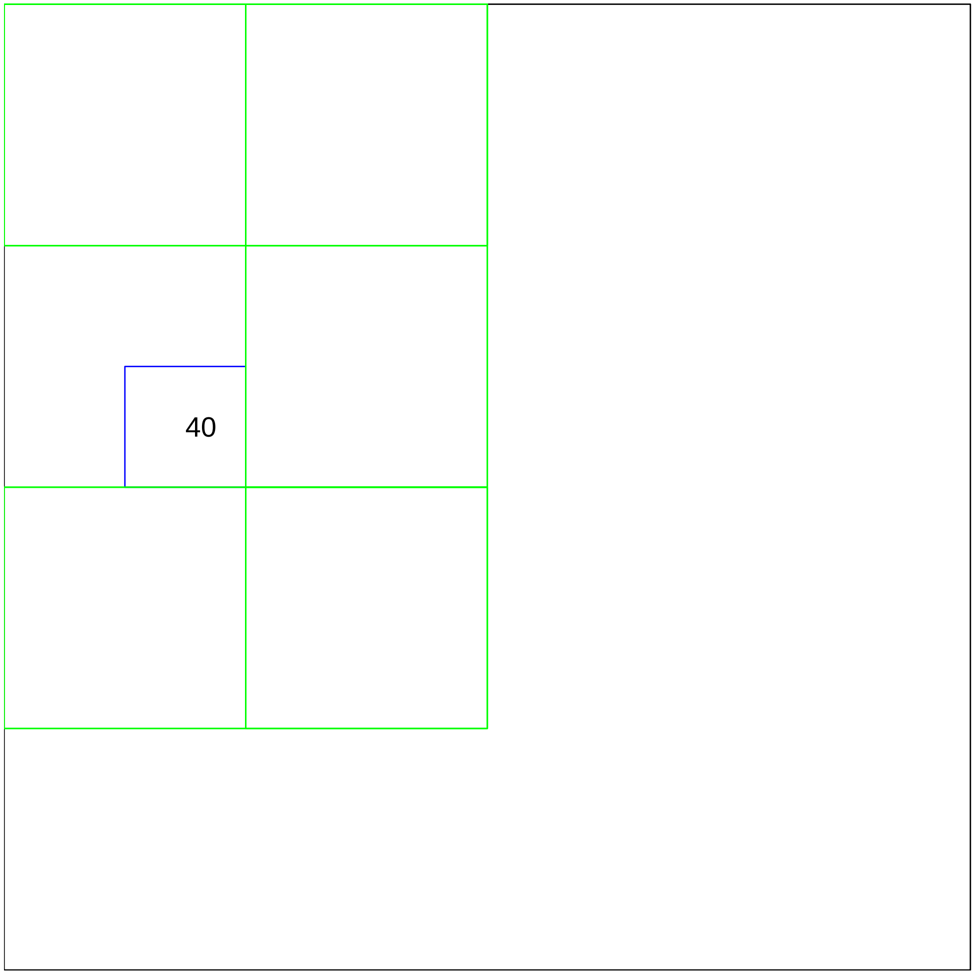

The green boxes are the neighbors of τ's parents.

Example: Consider box $\tau = 40$.

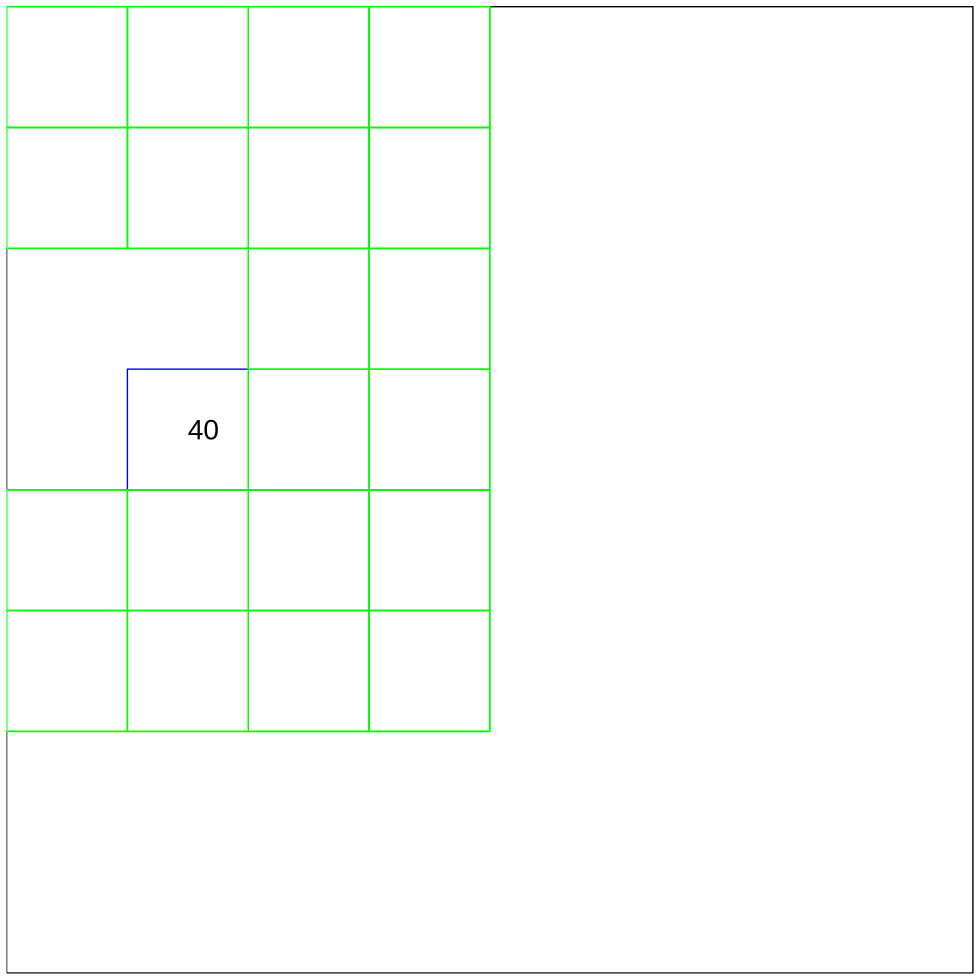

The green boxes are the children of the neighbors of τ's parents.

Example: Consider box $\tau = 40$.

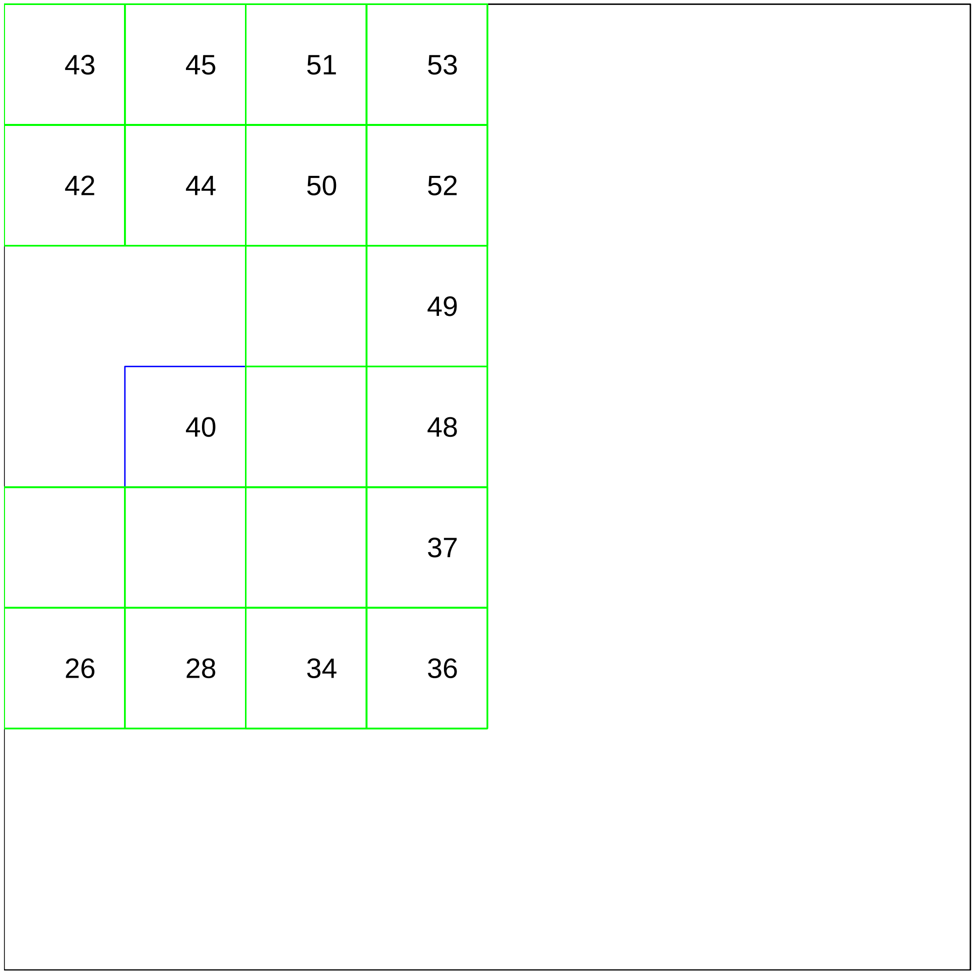

The green boxes are the children of the neighbors of τ's parents.

Only some of them are well-separated from τ.

We find: $\mathcal{L}_\tau^{(\mathrm{int})} = \{26, 28, 34, 36, 37, 42, 43, 44, 45, 48, 49, 50, 51, 52, 53\}$

Example: Consider box $\tau = 161$.

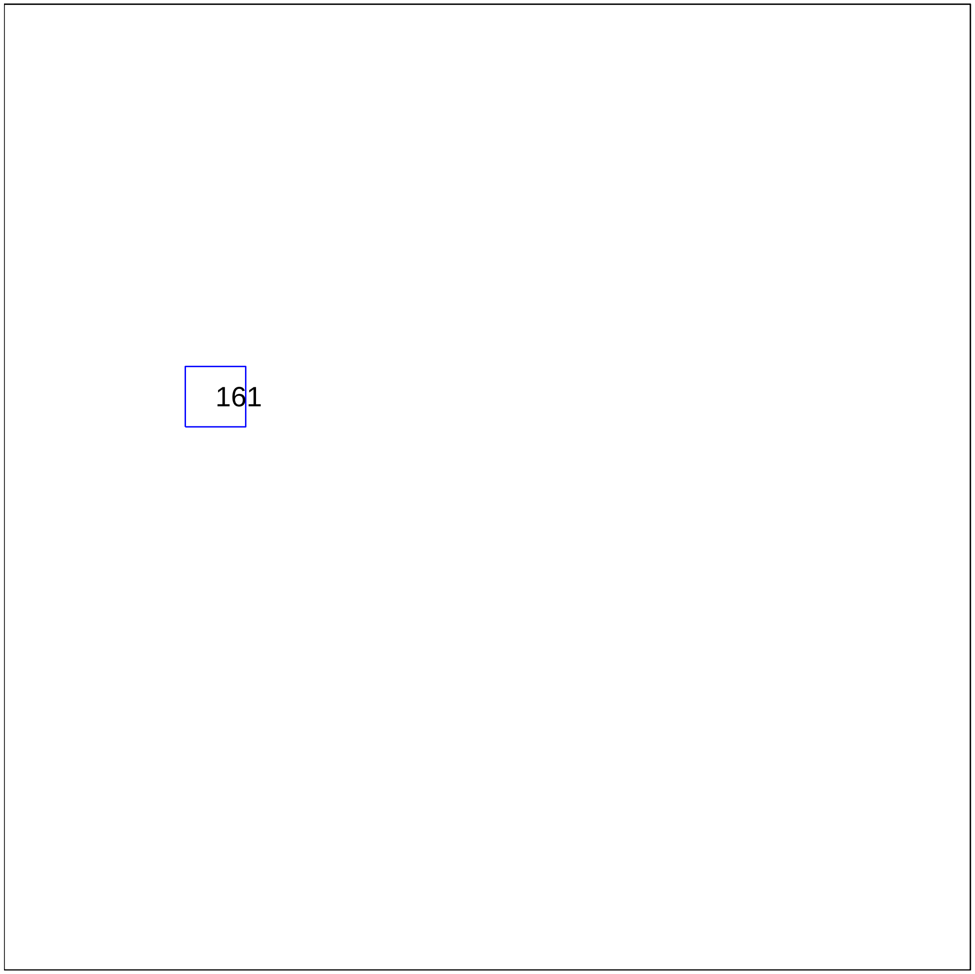

Example: Consider box $\tau = 161$.

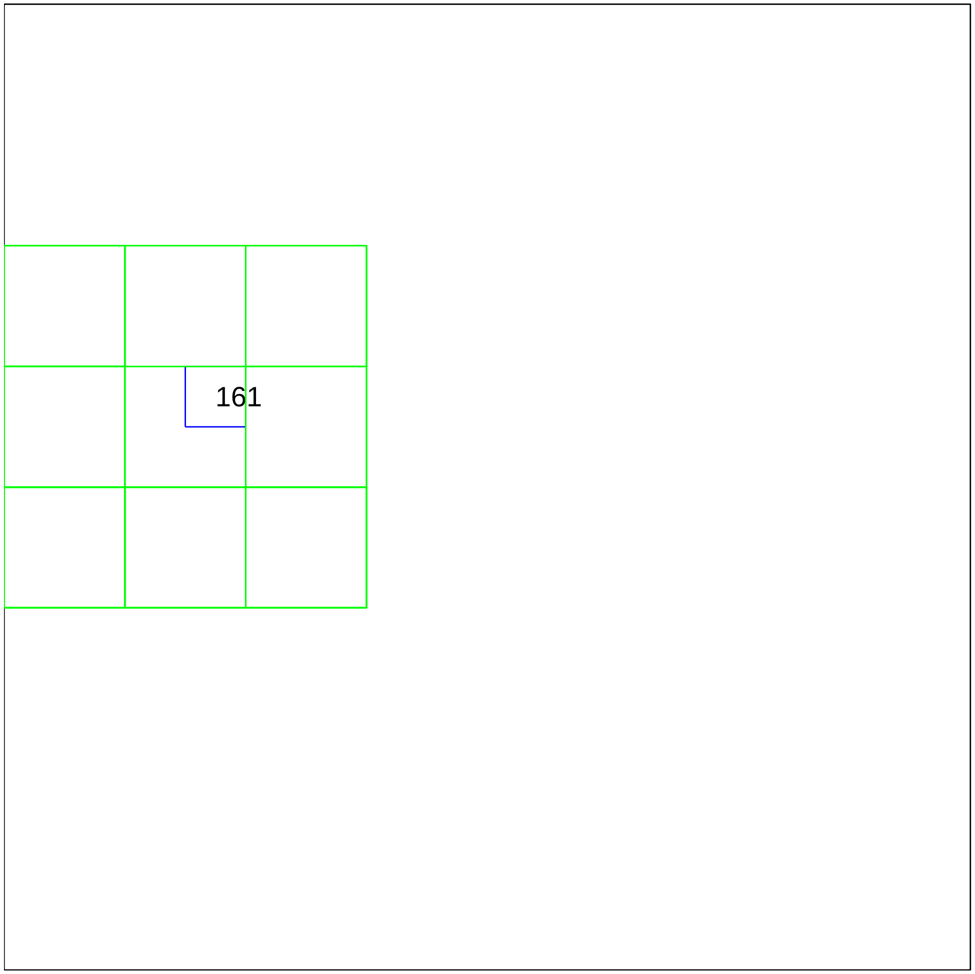

The green boxes are the neighbors of τ's parents.

Example: Consider box $\tau = 161$.

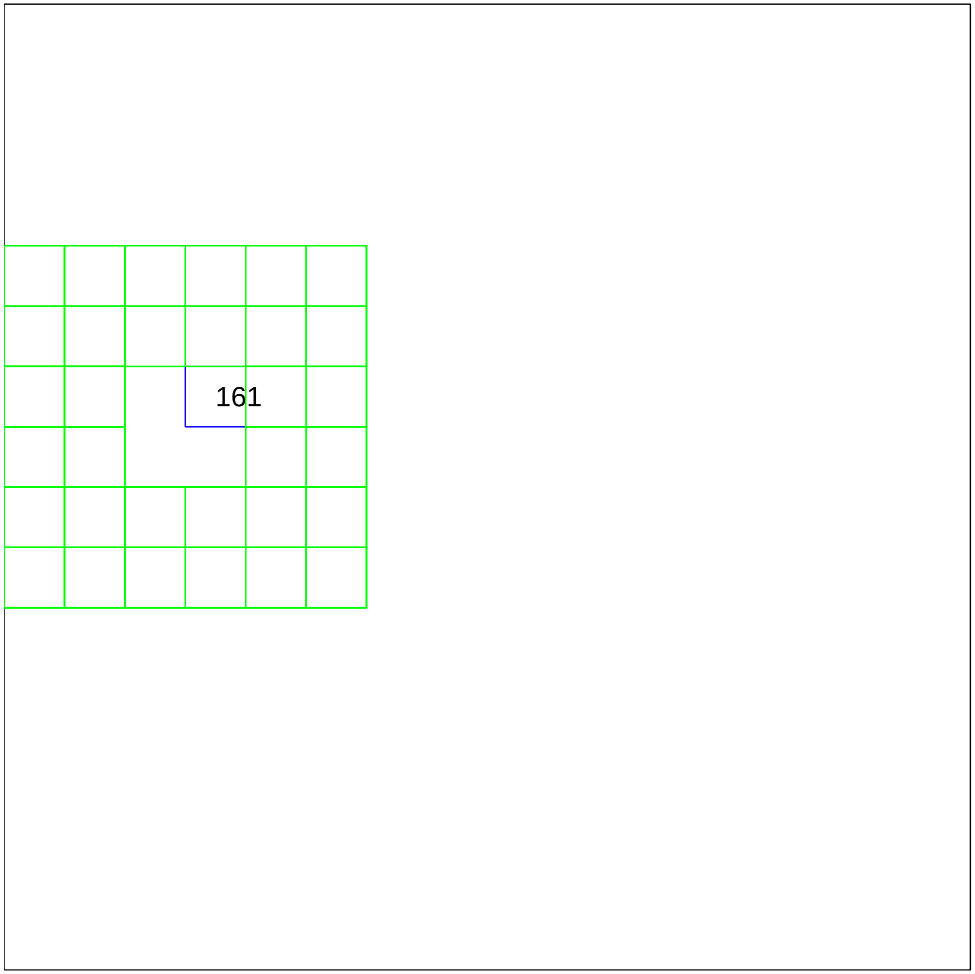

The green boxes are the children of the neighbors of τ's parents.

Example: Consider box $\tau = 161$.

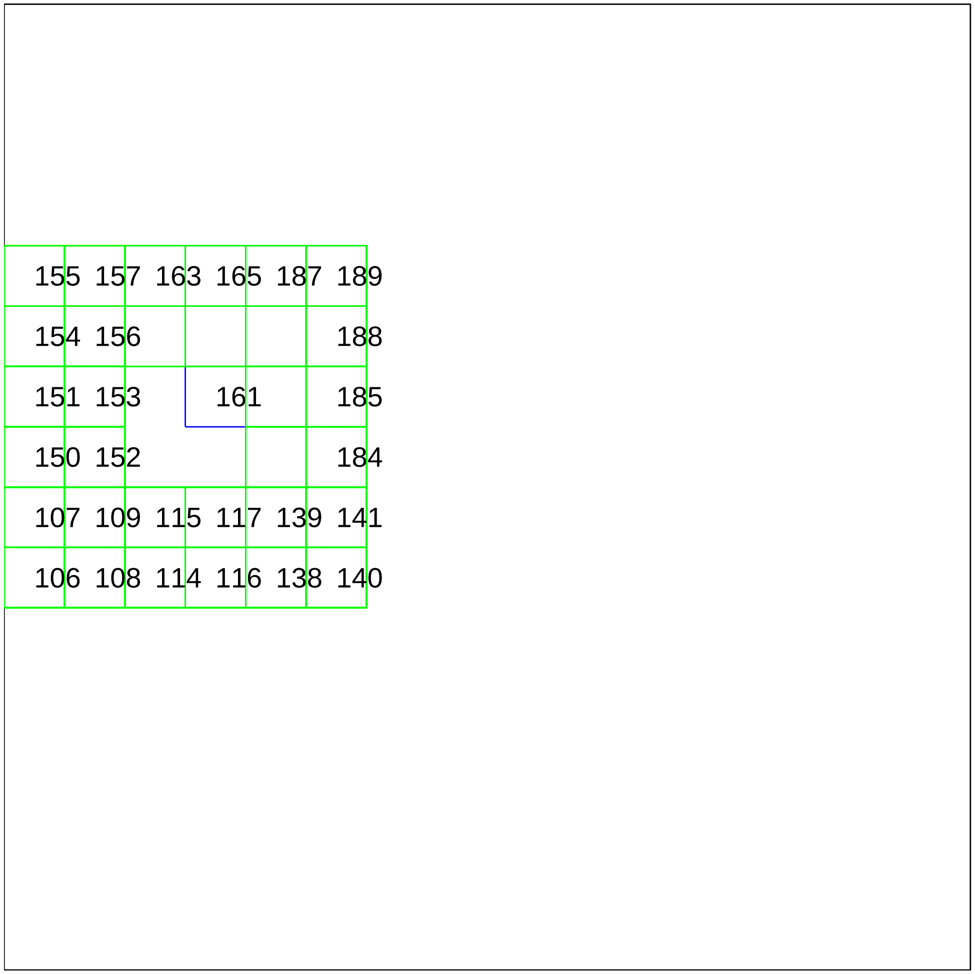

The green boxes are the children of the neighbors of τ's parents.

Only some of them are well-separated from τ.

We find: $\mathcal{L}_\tau^{(\mathrm{int})} = \{106, 107, 108, 109, 114, 115, 116, 117, 138, 139, 140, 141, 184, 185, 187, 188, 189, 150, 151, 152, 153, 154, 155, 156, 157, 163, 165\}$

Now let $\mathcal{L}_\tau^{(\mathrm{anc})}$ denote the list of all *ancestors* (parent, grand-parent, great-grand-parent, etc.). Then

$$\mathcal{L}_\tau^{(\mathrm{BH})} = \mathcal{L}_\tau^{(\mathrm{int})} \cup \bigcup_{\sigma \in \mathcal{L}_\tau^{(\mathrm{anc})}} \mathcal{L}_\sigma^{(\mathrm{int})}$$

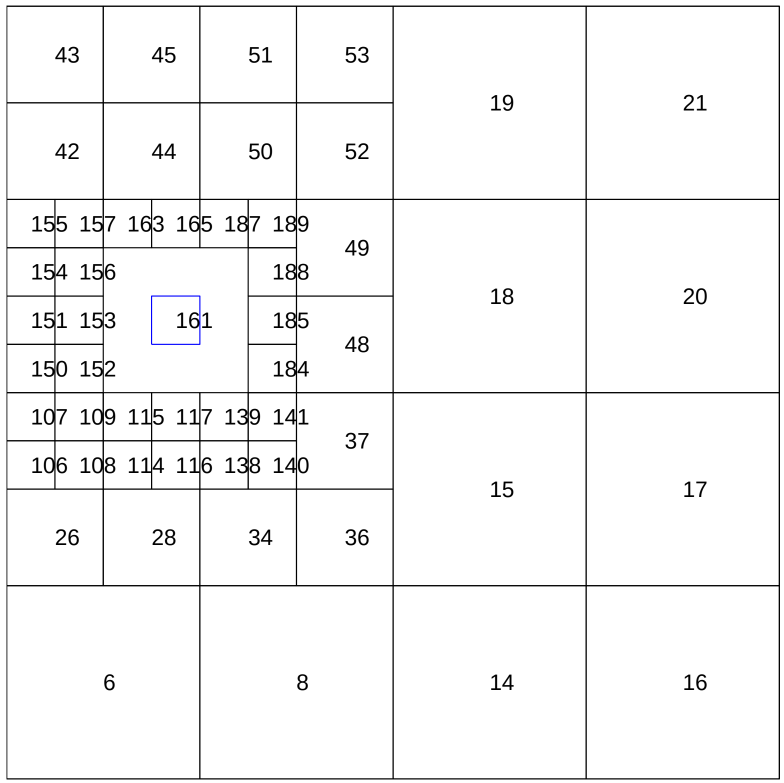
43
45
51
53
19
21
42
44
50
52
155
157
163
165
187
189
49
18
20
154
156
188
151
153
161
185
48
150
152
184
107
109
115
117
139
141
37
15
17
106
108
114
116
138
140
26
28
34
36
6
8
14
16

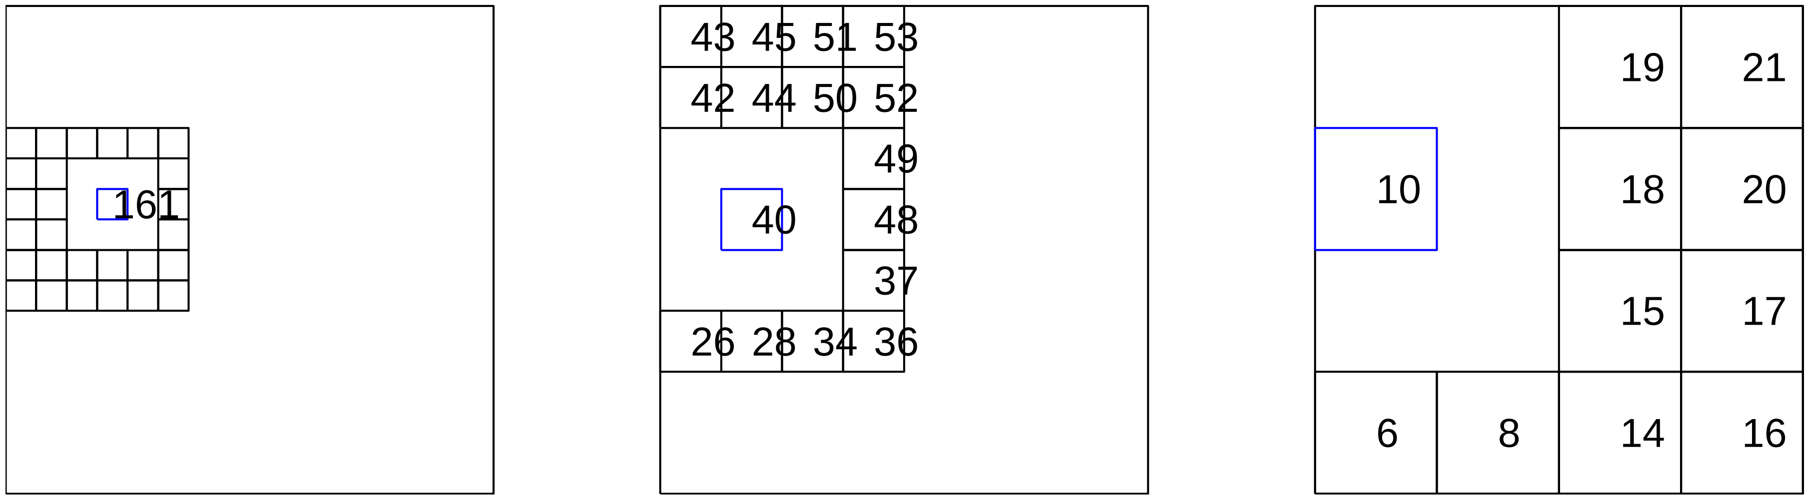

$$\mathcal{L}_{161}^{(\mathrm{anc})} = \{3,\ 10,\ 40\}$$

$$\begin{aligned}\mathcal{L}_{161}^{(\mathrm{BH})} &= \mathcal{L}_{3}^{(\mathrm{int})} \cup \mathcal{L}_{10}^{(\mathrm{int})} \cup \mathcal{L}_{40}^{(\mathrm{int})} \cup \mathcal{L}_{161}^{(\mathrm{int})} \\ &= \{6, 8, 14, 15, 16, 17, 18, 19, 20, 21, 26, 28, 34, 36, 37, 42, 43, 44, 45, 48, 49, \\ &\quad 50, 51, 52, 53, 106, 107, 108, 109, 114, 115, 116, 117, 138, 139, 140, 141, 150, \\ &\quad 151, 152, 153, 154, 155, 156, 157, 163, 165, 184, 185, 187, 188, 189\}\end{aligned}$$