# Communication Complexity of the Fast Multipole Method and its Algebraic Variants

Rio Yokota & David Keyes

KAUST
King Abdullah University of Science and Technology

2014 CBMS-NSF Conference:
Fast Direct Solvers for Elliptic PDEs
June 23–29, 2014, Dartmouth College

## Computational Complexity $\mathcal{O}(N^2) \rightarrow \mathcal{O}(N)$

Given a problem of size *N*, how many arithmetic operations does the algorithm perform?

## Communication Complexity $\mathcal{O}(P) \rightarrow \mathcal{O}(\log P)$

Given a problem on *P* parallel processes, how much data does the algorithm have to send?

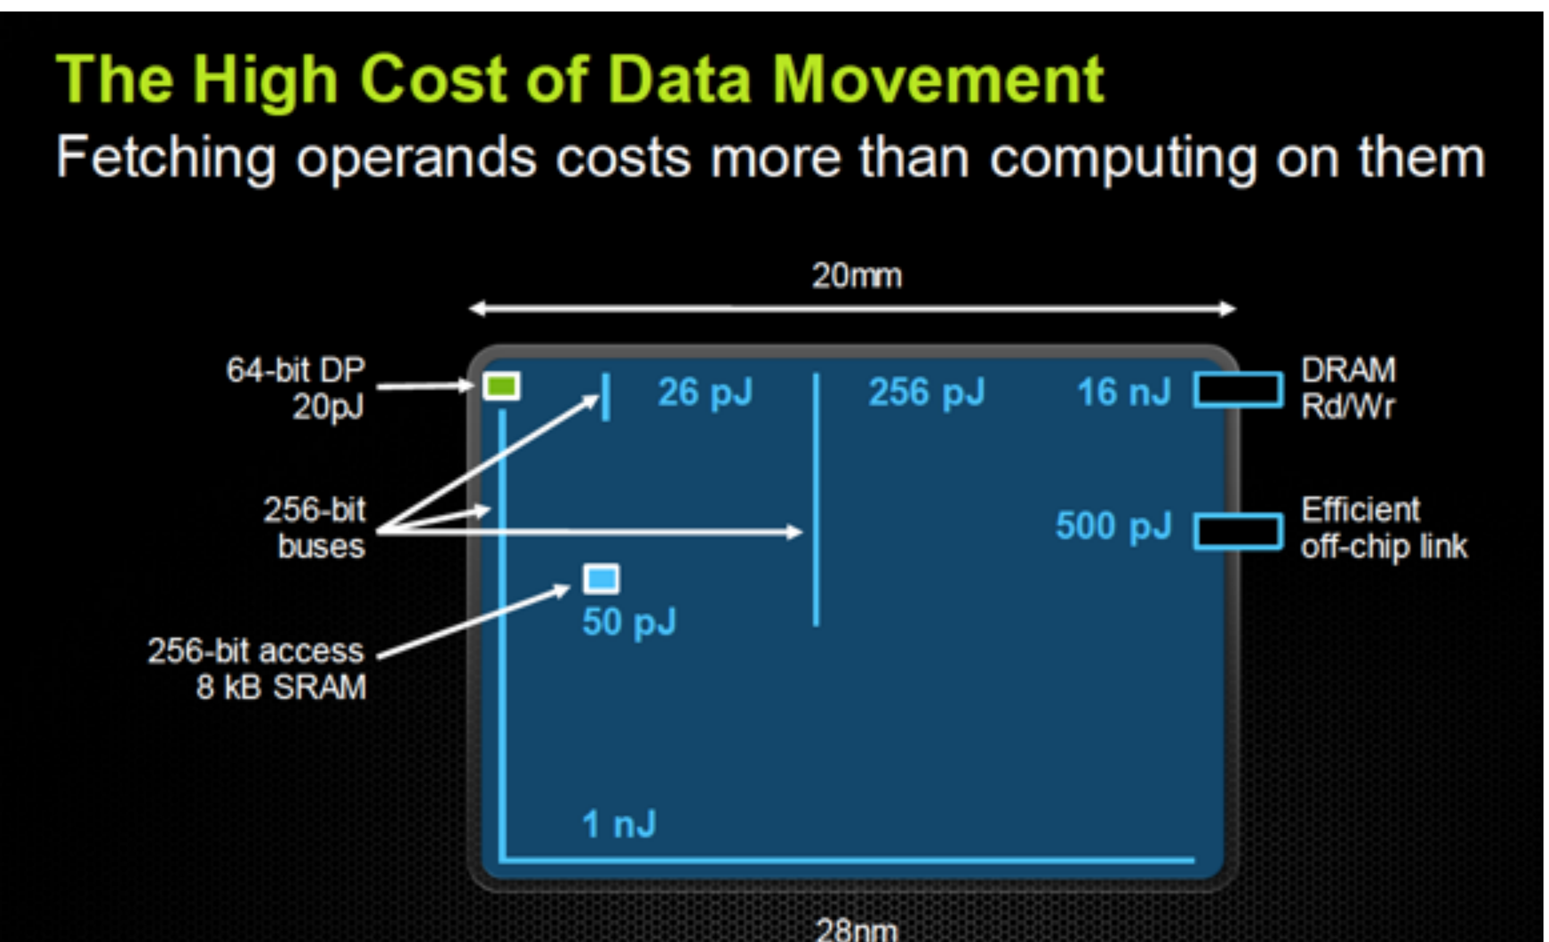


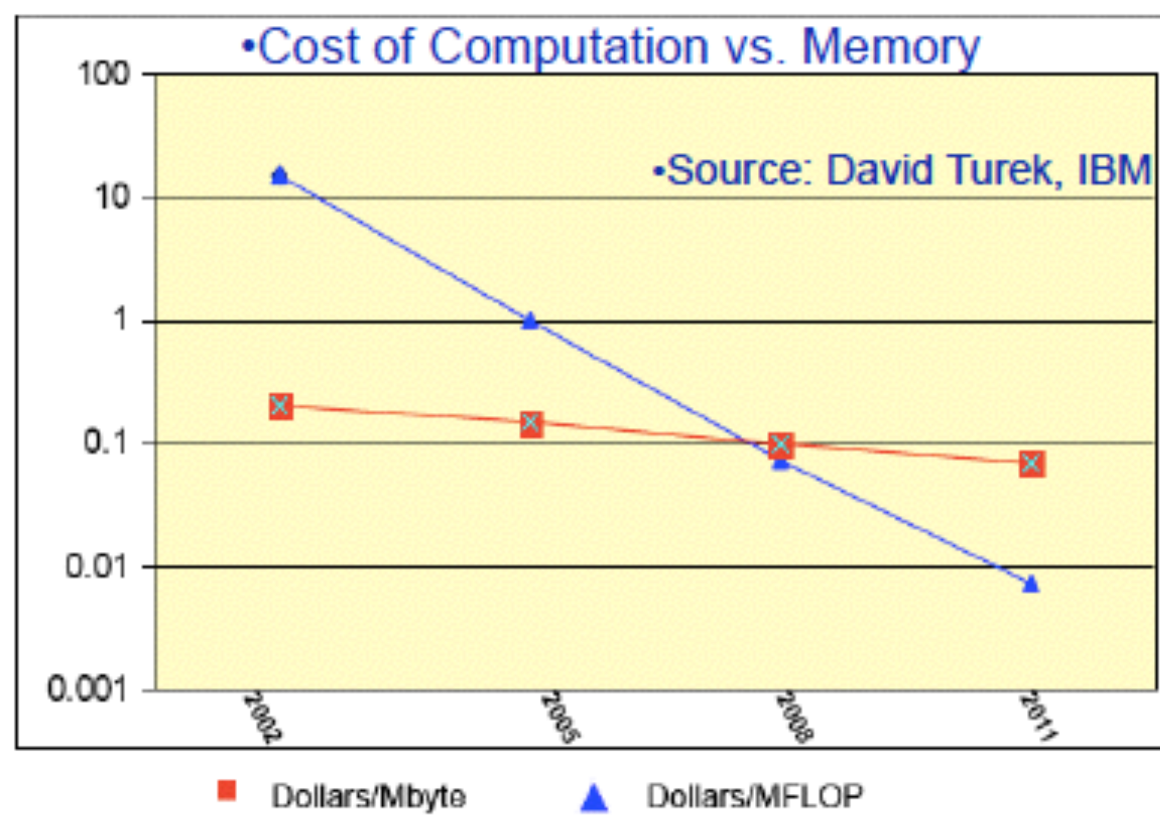


Source: Bill Dally, NVIDIA

# Communication Complexity of FMM

SIAM J. Sci. Comput.
Vol. 19, No. 2, pp. 635–656, March 1998


**PROVABLY GOOD PARTITIONING AND LOAD BALANCING ALGORITHMS FOR PARALLEL ADAPTIVE N-BODY SIMULATION***

SHANG-HUA TENG†

A massively parallel adaptive fast-multipole method on heterogeneous architectures

Ilya Lashuk*, Aparna Chandramowlishwaran*, Harper Langston*, Tuan-Anh Nguyen*, Rahul Sampath*, Aashay Shringarpure*, Richard Vuduc*, Lexing Ying†, Denis Zorin‡, and George Biros*
* Georgia Institute of Technology, Atlanta, GA 30332
† University of Texas at Austin, Austin, TX 78712
‡ New York University, New York, NY 10002
ilashuk@cc.gatech.edu, aparna@cc.gatech.edu, harper@cc.gatech.edu,
tuananh.nguyen@gatech.edu rahul.sampath@gmail.com, aashay.shringarpure@gmail.com,
richie@cc.gatech.edu, lexing@math.utexas.edu, dzorin@cs.nyu.edu, gbiros@acm.org

$$\sum_{i}^{\log P-1} 2^i = \mathcal{O}(P)$$

$$\sum_{i}^{\log P-1} min(2^{\log P-i-1}, 2^i) = \mathcal{O}(\sqrt{P})$$

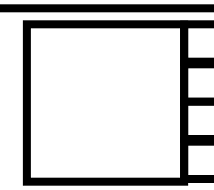

| Reference | Processes | | Data per Process | | Communication complexity |
|---|---|---|---|---|---|
| Teng | $\mathcal{O}(P)$ | | $\mathcal{O}\left((N/P)^{2/3}(\log N+\mu)^{1/3}\right)$ | | $\mathcal{O}\left(P(N/P)^{2/3}(\log N+\mu)^{1/3}\right)$ |
| Lashuk *et al.* | $\mathcal{O}(\sqrt{P})$ | | $\mathcal{O}\left((N/P)^{2/3}\right)$ | | $\mathcal{O}\left(\sqrt{P}(N/P)^{2/3}\right)$ |
| Ibeid *et al.* | Global | Local | Global | Local | Global + Local |
| | $\mathcal{O}(\log P)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}\left((N/P)^{2/3}\right)$ | $\mathcal{O}\left(\log P+(N/P)^{2/3}\right)$ |

# Uniform Case

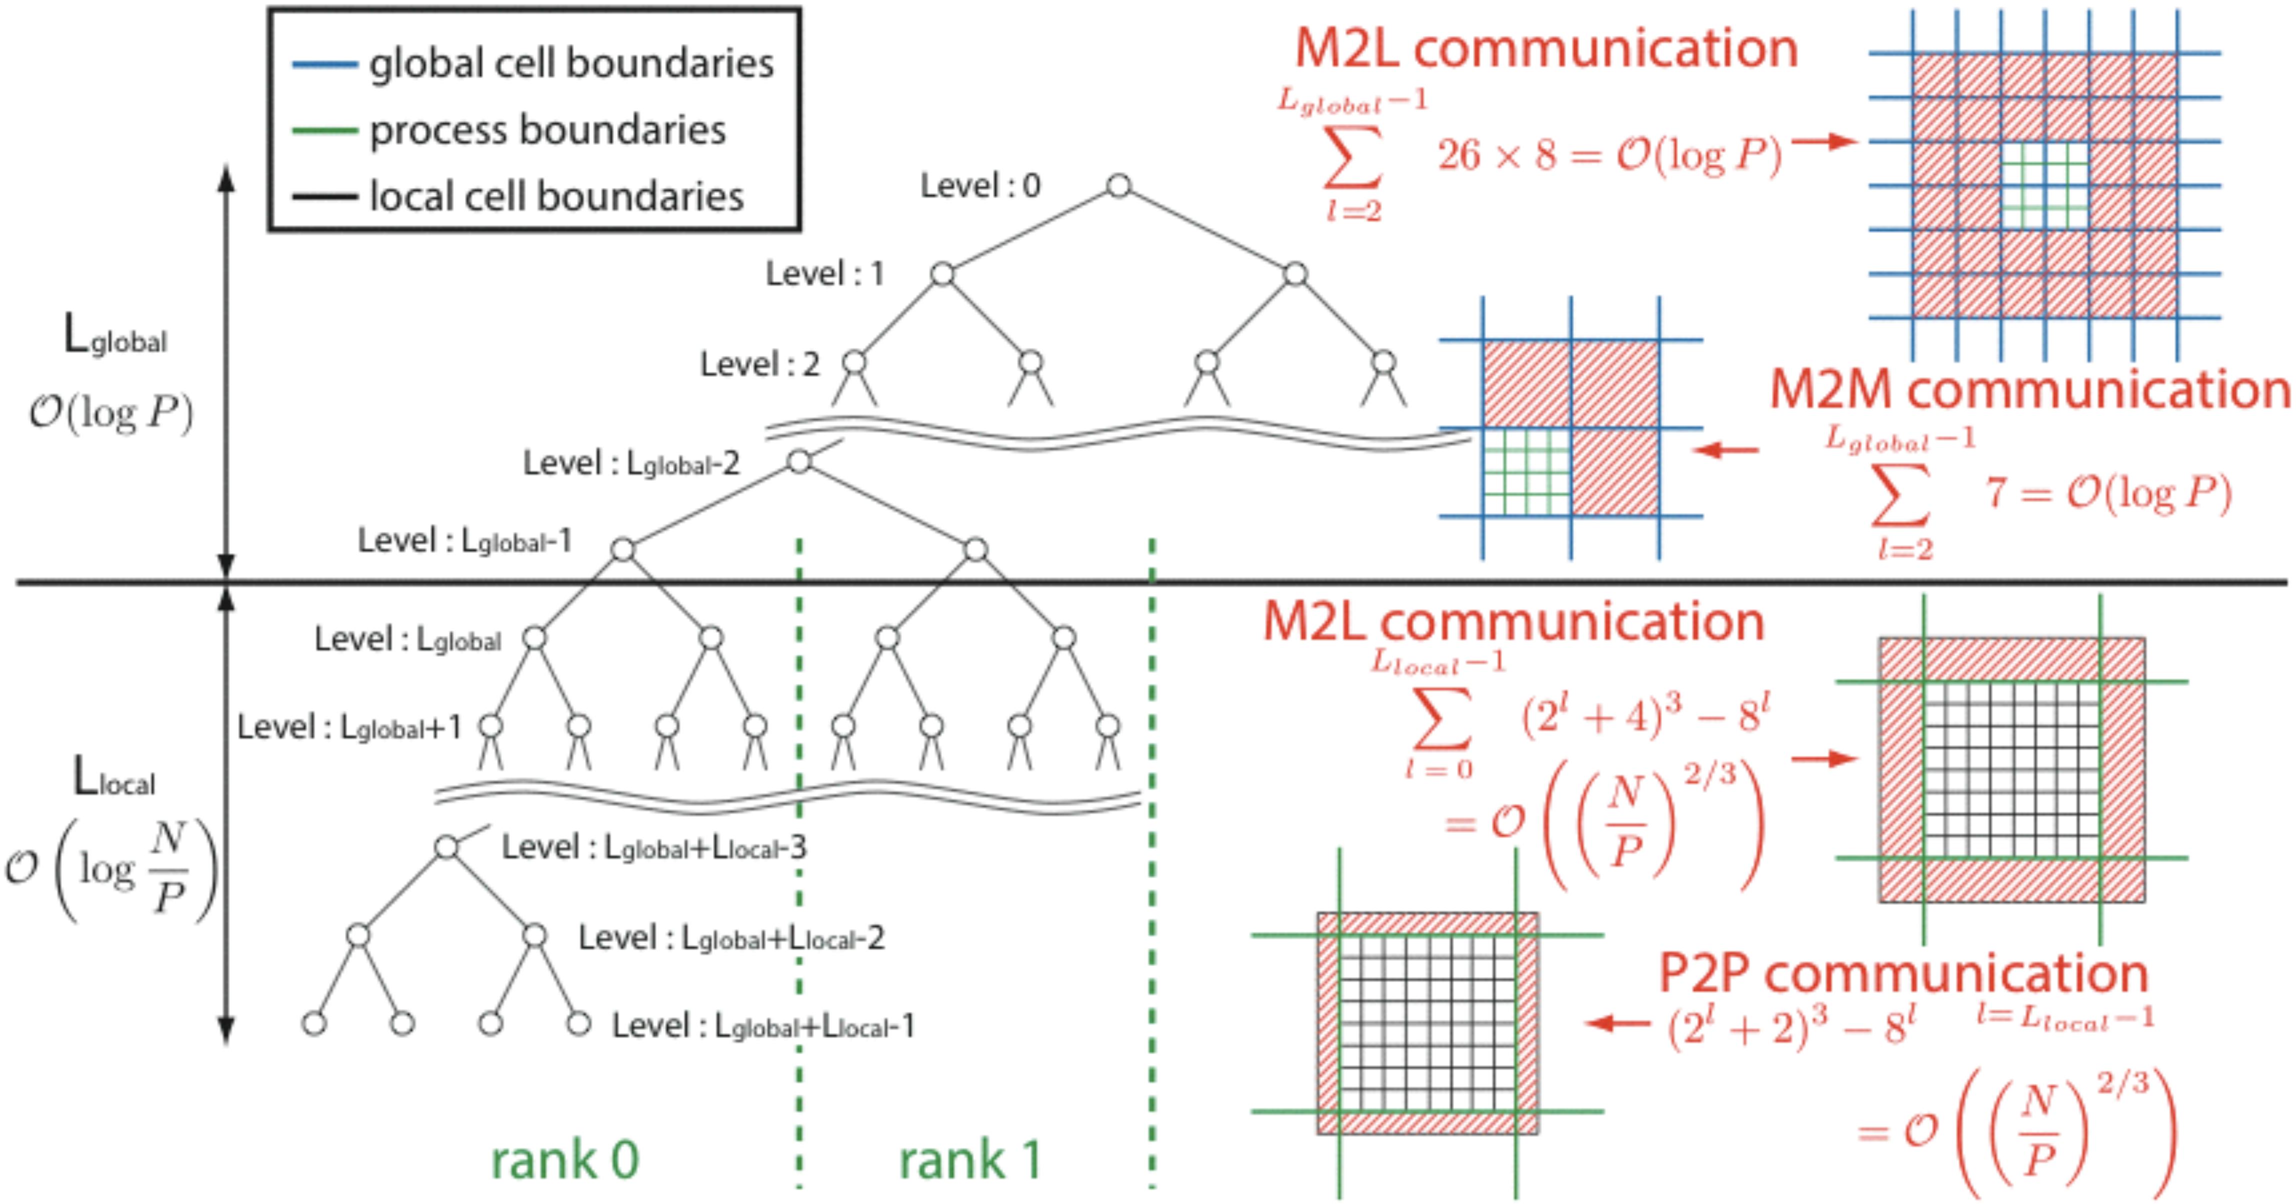


| | Processes | Cells per level | Cells per Process | Communication |
|---|---|---|---|---|
| Global M2L | $\sum_i^{\log P} 26$ | $26 \times 8$ | 8 | $\mathcal{O}(\log P)$ |
| Global M2M | $\sum_i^{\log P} 7$ | 7 | 1 | $\mathcal{O}(\log P)$ |
| Local M2L | 26 | $(2^i+4)^3 - 8^i$ | $\sum_i^{\log_8(N/P)} (2^i+4)^3 - 8^i$ | $\mathcal{O}((N/P)^{2/3})$ |
| Local P2P | 26 | $(2^i+2)^3 - 8^i$ | $(2^{\log_8(N/P)}+2)^3 - 8^{\log_8(N/P)}$ | $\mathcal{O}((N/P)^{2/3})$ |

# Nonuniform Case

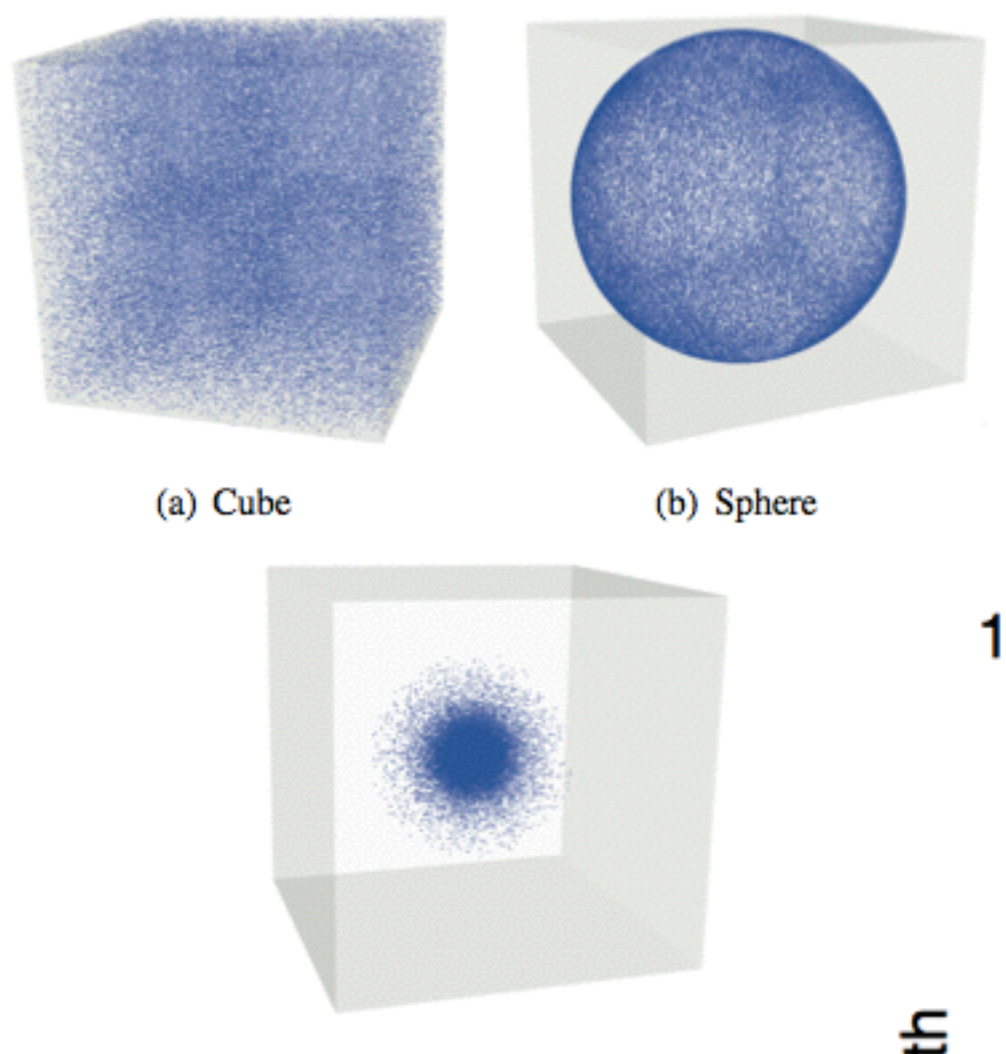

(a) Cube (b) Sphere

(c) Plummer

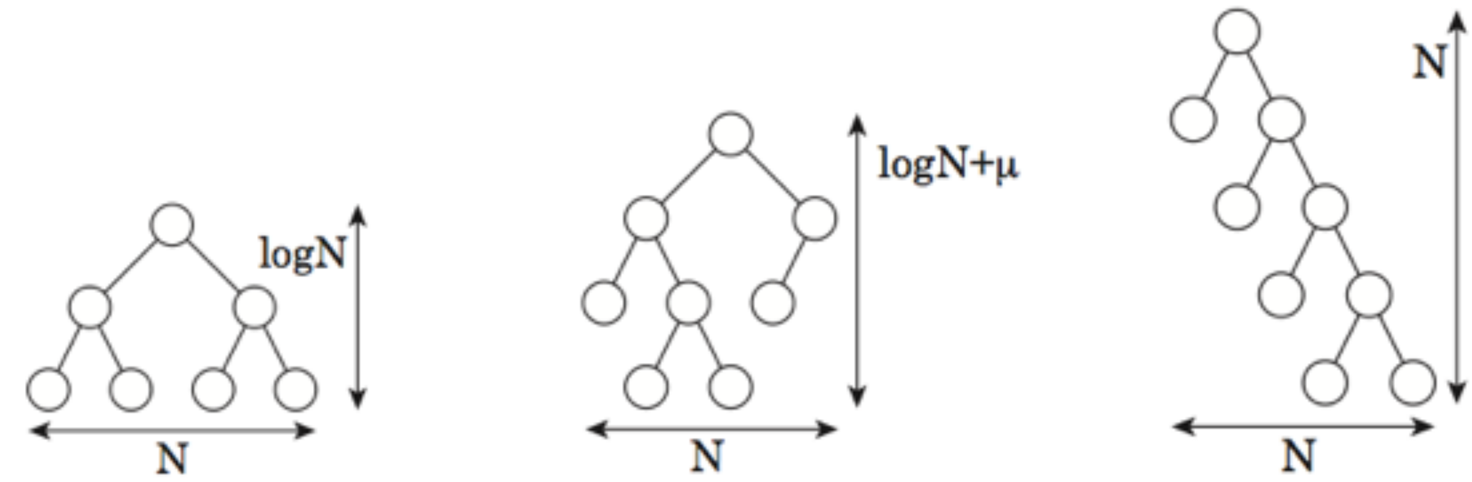


(a) Uniform (b) Nonuniform (c) Pathological

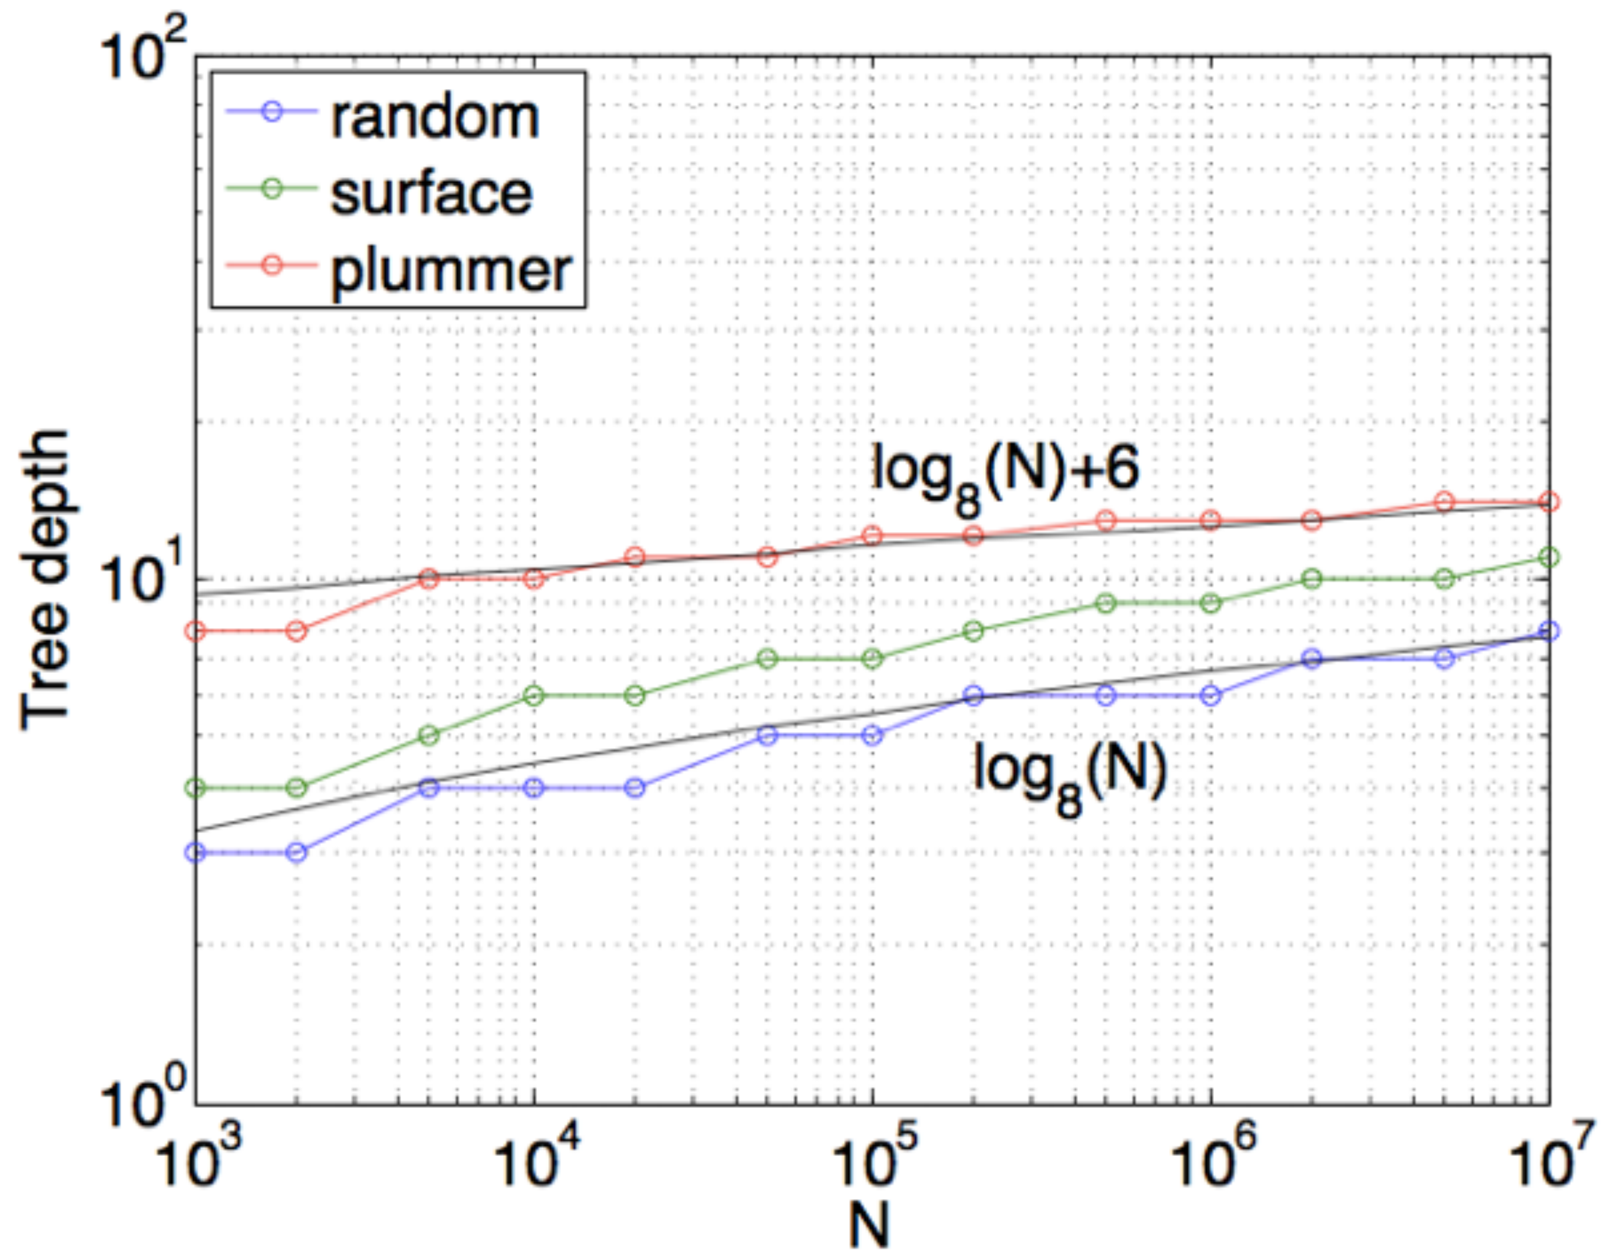

# Nonuniform Case

| | Processes | Cells per level | Cells per Process | Communication |
|---|---|---|---|---|
| Global M2L | $\sum_i^{\log P} 26$ | $26 \times 8$ | $8$ | $\mathcal{O}(\log P)$ |
| Global M2M | $\sum_i^{\log P} 7$ | $7$ | $1$ | $\mathcal{O}(\log P)$ |
| Local M2L | $26$ | $(2^i+4)^3 - 8^i$ | $\sum_i^{\log_8(N/P)} (2^i+4)^3 - 8^i$ | $\mathcal{O}((N/P)^{2/3})$ |
| Local P2P | $26$ | $(2^i+2)^3 - 8^i$ | $(2^{\log_8(N/P)}+2)^3 - 8^{\log_8(N/P)}$ | $\mathcal{O}((N/P)^{2/3})$ |

↓

| | Processes | Cells per level | Cells per Process | Communication |
|---|---|---|---|---|
| Global M2L | $\sum_i^{\log P} \mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(\log P)$ |
| Global M2M | $\sum_i^{\log P} \mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(\log P)$ |
| Local M2L | $\mathcal{O}(1)$ | $\mathcal{O}(4^i)$ | $\sum_i^{\log_8(N/P)} \mathcal{O}(4^i)$ | $\mathcal{O}((N/P)^{2/3})$ |
| Local P2P | $\mathcal{O}(1)$ | $\mathcal{O}(4^i)$ | $\mathcal{O}(4^{\log_8(N/P)})$ | $\mathcal{O}((N/P)^{2/3})$ |

# Algebraic Case (Mat-Vec)

$$y = \left( \sum_{(i,j) \in D} A_{ij} \right) x + \left( \sum_{(i,j) \in L} U_i S_{ij} V_j^t \right) x = \underbrace{\sum_{(i,j) \in D} A_{ij} x_j}_{\text{Dense mat-vecs operations}} + \underbrace{\sum_{i \in I} U_i \underbrace{\sum_{(i,j) \in L} S_{ij} \underbrace{V_j^t x}_{\text{Upsweep}}}_{\text{Coupling phase}}}_{\text{Downsweep}}$$

| | Processes | Cells per level | Cells per Process | Communication |
|---|---|---|---|---|
| Global M2L | $\sum_i^{\log P} \mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(\log P)$ |
| Global M2M | $\sum_i^{\log P} \mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(\log P)$ |
| Local M2L | $\mathcal{O}(1)$ | $\mathcal{O}(4^i)$ | $\sum_i^{\log_8(N/P)} \mathcal{O}(4^i)$ | $\mathcal{O}((N/P)^{2/3})$ |
| Local P2P | $\mathcal{O}(1)$ | $\mathcal{O}(4^i)$ | $\mathcal{O}(4^{\log_8(N/P)})$ | $\mathcal{O}((N/P)^{2/3})$ |

| Matrix Operation | FMM operation | Processes | Blocks per level | Blocks per Process | Communication |
|---|---|---|---|---|---|
| Global $\sum S_{ij} \hat{x}_j$ | Global M2L | $\sum_i^{\log P} \mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(\log P)$ |
| Global $\sum F_k^t \hat{x}_k$ | Global M2M | $\sum_i^{\log P} \mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}(\log P)$ |
| Local $\sum S_{ij} \hat{x}_j$ | Local M2L | $\mathcal{O}(1)$ | $\mathcal{O}(2^{(d-1)i})$ | $\sum_i^{\log_{2^d}(N/P)} \mathcal{O}(2^{(d-1)i})$ | $\mathcal{O}((N/P)^{\frac{d-1}{d}})$ |
| Local $\sum A_{ij} x_j$ | Local P2P | $\mathcal{O}(1)$ | $\mathcal{O}(2^{(d-1)i})$ | $\mathcal{O}(2^{(d-1)\log_{2^d}(N/P)})$ | $\mathcal{O}((N/P)^{\frac{d-1}{d}})$ |

# Asymptotic Constants

## FMM

| Type of expansion (+M2L acceleration) | Storage | Arithmetic |
|---|---|---|
| Cartesian Taylor | $\mathcal{O}(p^3)$ | $\mathcal{O}(p^6)$ |
| Cartesian Chebychev | $\mathcal{O}(p^3)$ | $\mathcal{O}(p^6)$ |
| Spherial harmonics | $\mathcal{O}(p^2)$ | $\mathcal{O}(p^4)$ |
| Spherial harmonics+rotation | $\mathcal{O}(p^2)$ | $\mathcal{O}(p^3)$ |
| Spherial harmonics+FFT | $\mathcal{O}(p^2)$ | $\mathcal{O}(p^2 \log^2 p)$ |
| Planewave | $\mathcal{O}(p^2)$ | $\mathcal{O}(p^3)$ |
| Equivalent charges | $\mathcal{O}(p^2)$ | $\mathcal{O}(p^4)$ |
| Equivalent charges+FFT | $\mathcal{O}(p^3)$ | $\mathcal{O}(p^3 \log p)$ |

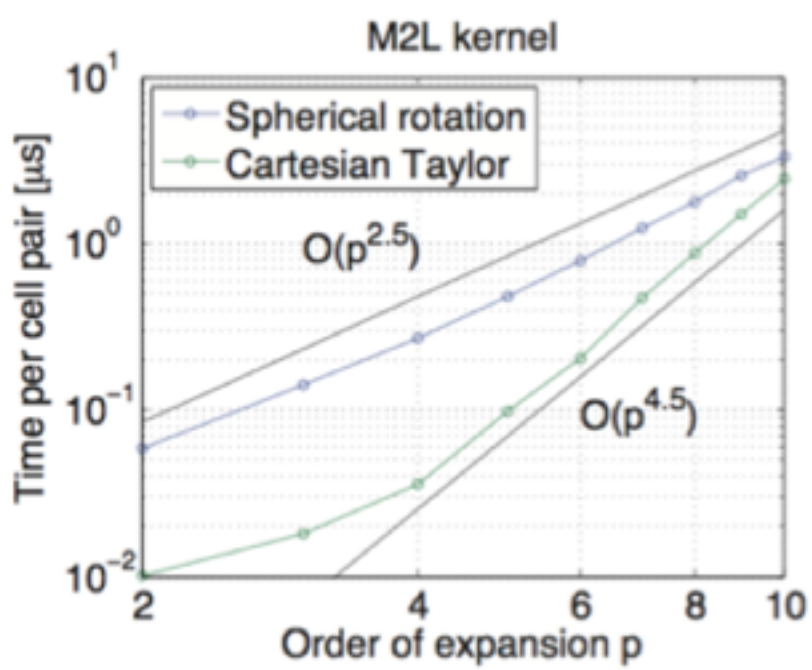


$$Better\ complexity \neq Better\ performance$$

## H-matrix

| Type of low rank approximation | Reference |
|---|---|
| Rank-revealing LU | Pan (2000) |
| Rank-revealing QR | Gu & Eisenstat (1996) |
| Pivoted QR | Kong et al. (2011) |
| Truncated SVD | Grasedyck & Hackbusch (2003) |
| Randomized SVD | Liberty et al. (2007) |
| Adaptive cross approximation | Rjasanow (2002) |
| Hybrid cross approximation | Börm (2005) |
| Chebychev interpolation | Dutt et al. (1996) |

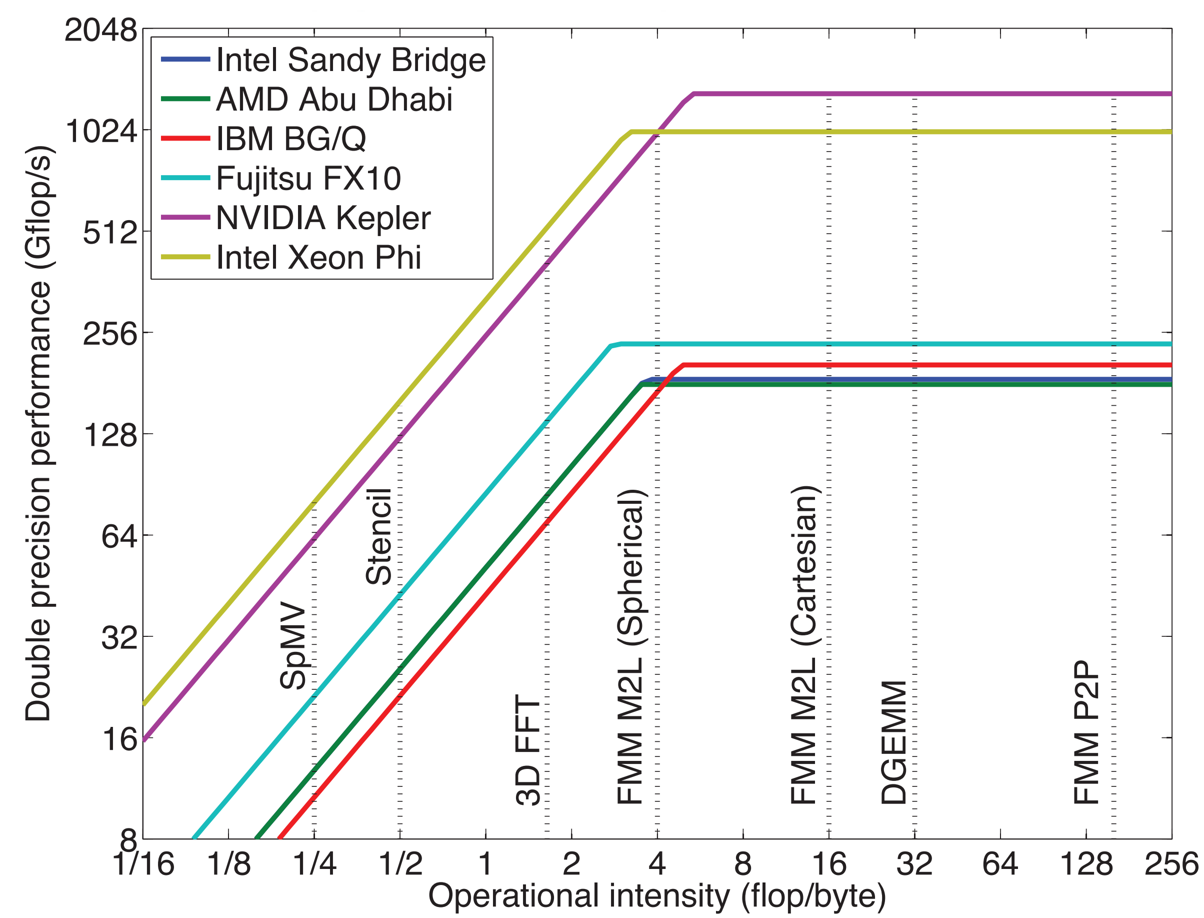

# Finding Admissible Blocks

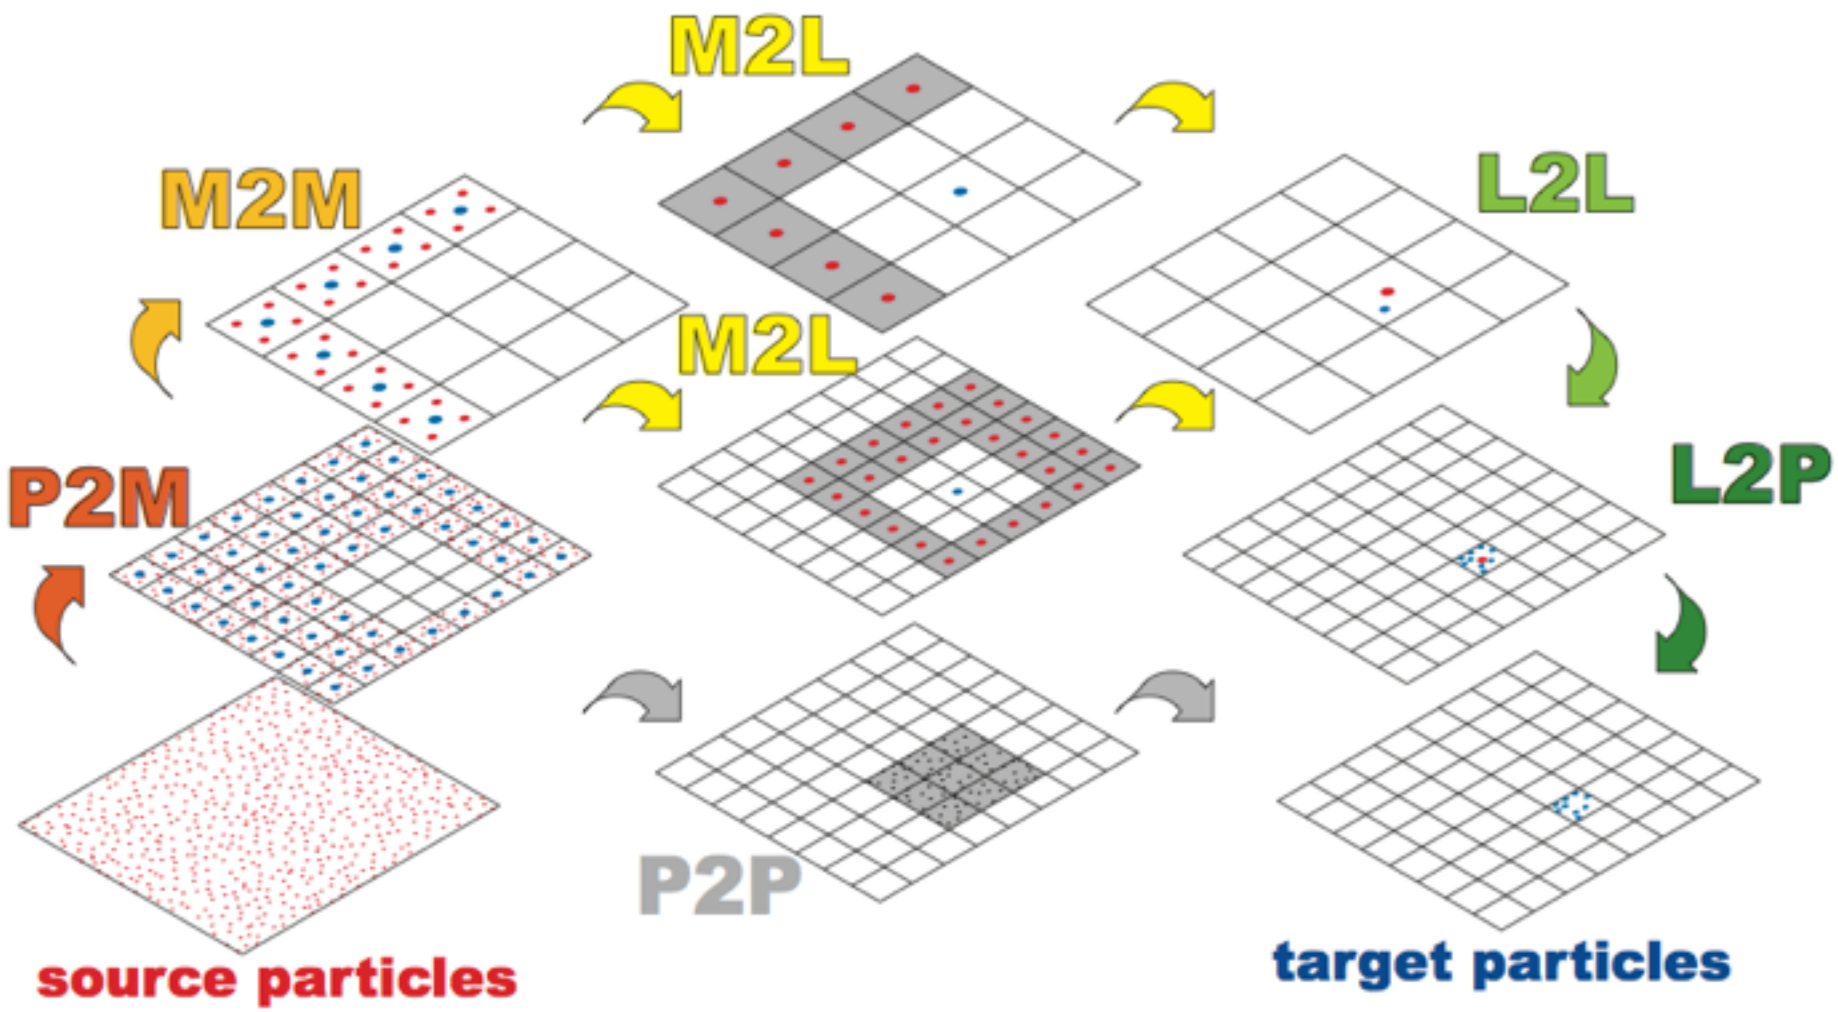


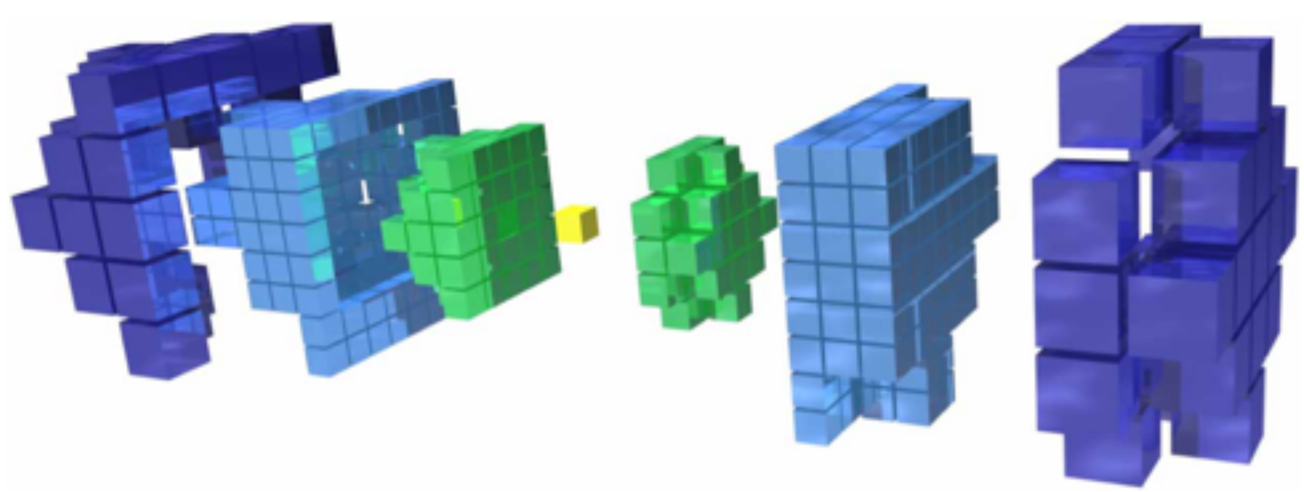

Neighbor List Per Level Per Block?

# Dual Tree Traversal

W. Dehnen, J. Comput. Phys. 179; 27-42 (2002)

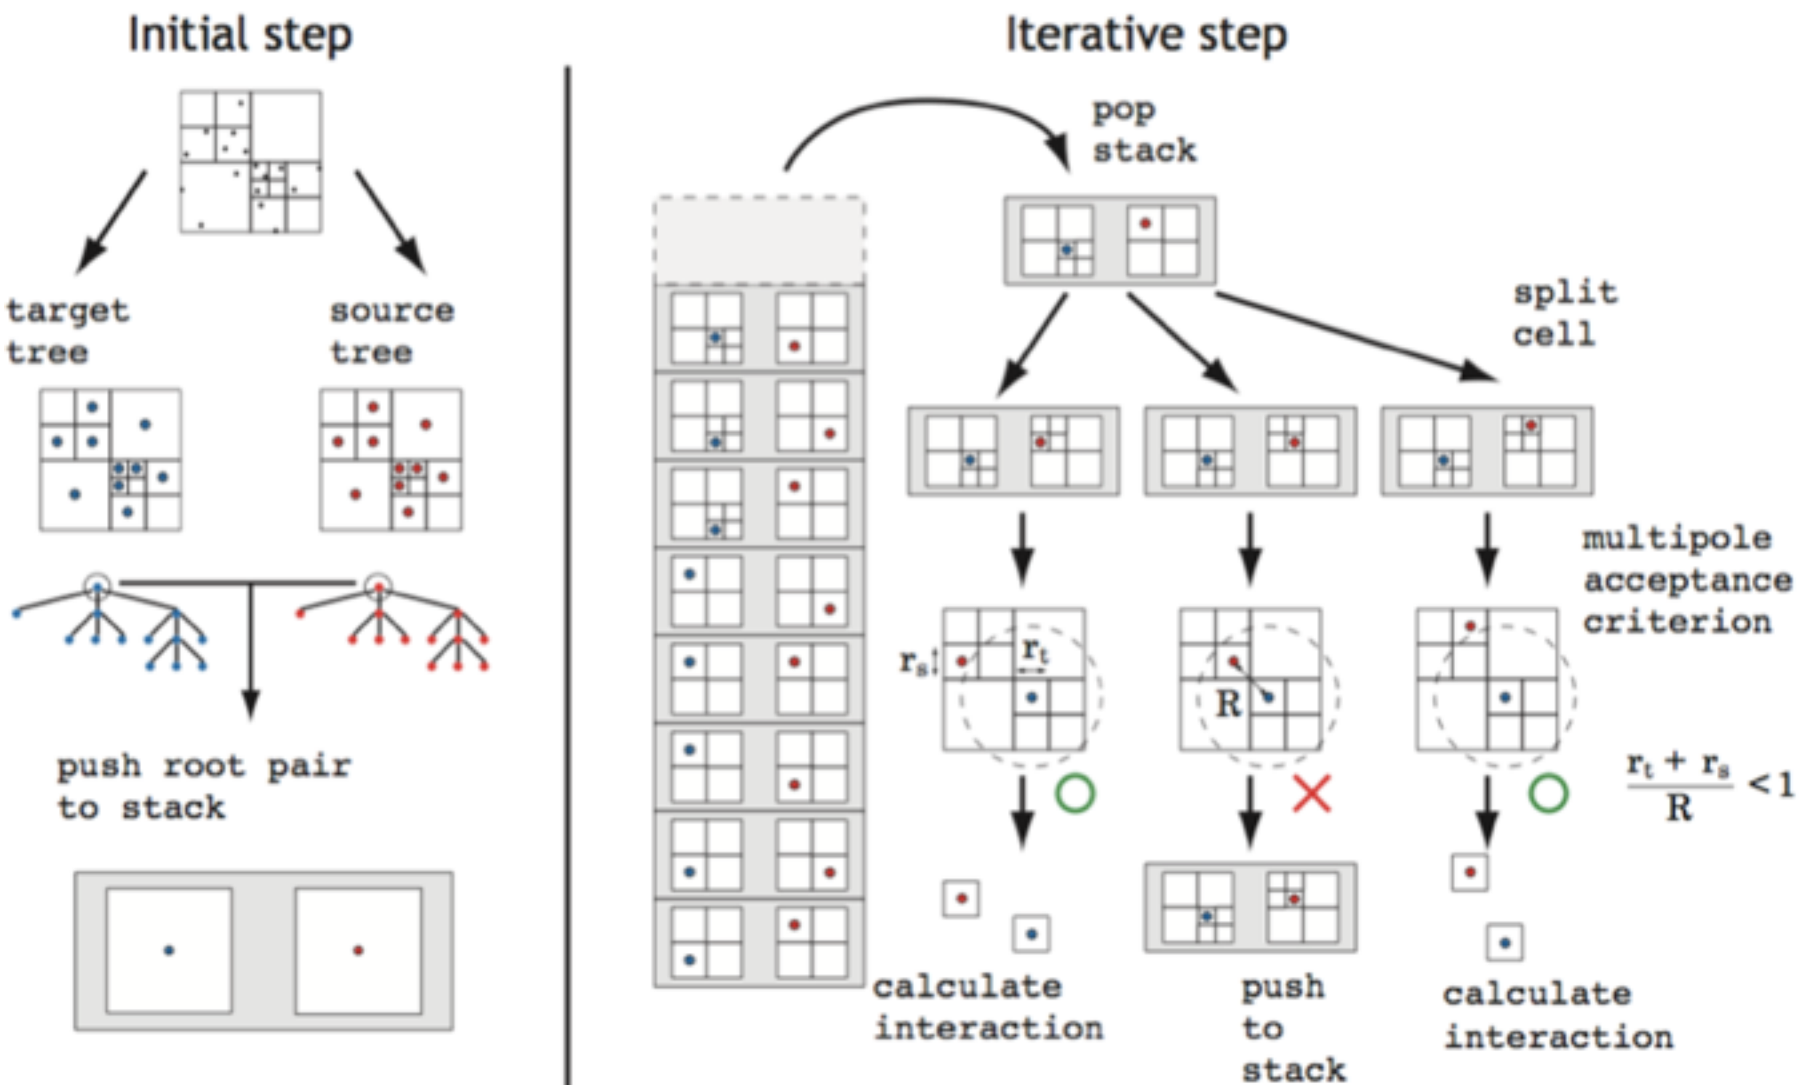


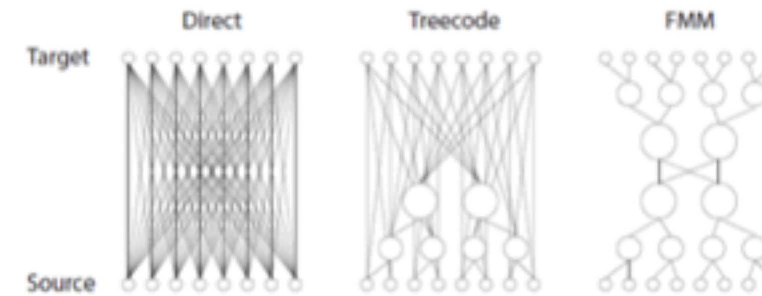



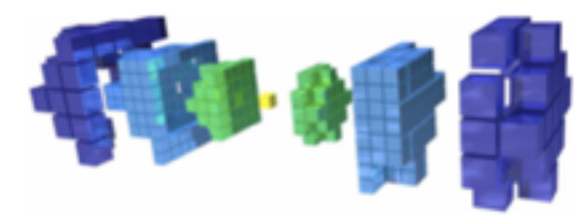

No Neighbor List Needed

S.-H. Teng, SIAM J. Sci. Comput. 19; 635-656 (1998)

# Cells don't have to be cubic

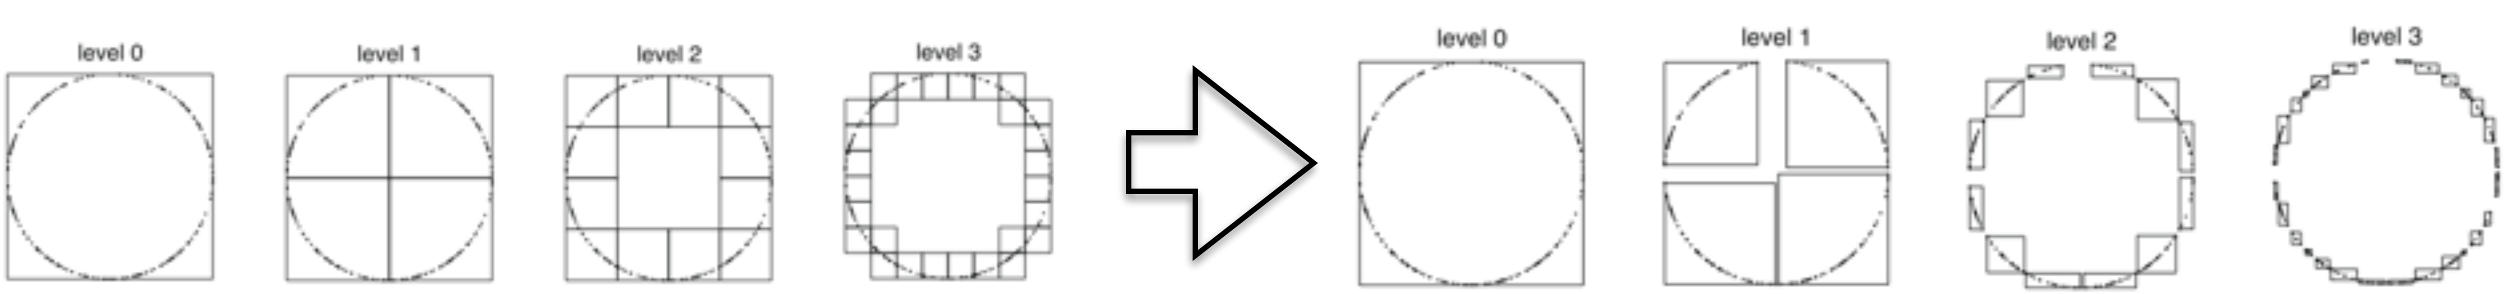

# Dual Tree Traversal

- No need to explicitly form interaction lists
- The definition of well-separatedness (size of neighbor region) can be adjusted flexibly without modifying the code
- The resulting neighbor region and M2L interaction region naturally have a spherical shape (as opposed to the suboptimal cubic shape)
- It is applicable to adaptive trees without any modification
- It lends itself to MPI parallelization without any modification (by simply using the local essential tree as the source tree)
- It is easily extendable to periodic boundary conditions (by traversing all periodic image trees as the source tree)
- It can handle mutual M2L interaction, and can satisfy Newtons third law (M2L is neither target centric nor source centric, but completely symmetric)
- It works well with task based threading tools like Intel TBB, Cilk, MassiveThreads, etc., where tasks are spawned while the tree is traversed
- The cells don't have to be cubic. For example, high aspect ratio rectangles or hierarchical K-means is permitted.
- It can be implemented in less than 100 lines of code, and is therefore trivial to debug

# h-p FMM

$$\theta = \frac{|\mathbf{x}_j - \mathbf{x}_*|}{|\mathbf{x}_i - \mathbf{x}_*|} < 1$$

Error bound : $O(\theta^p)$

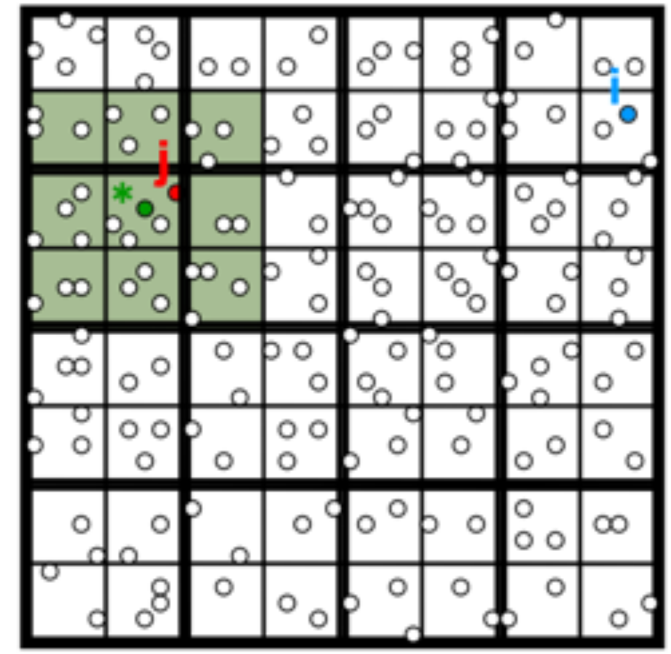


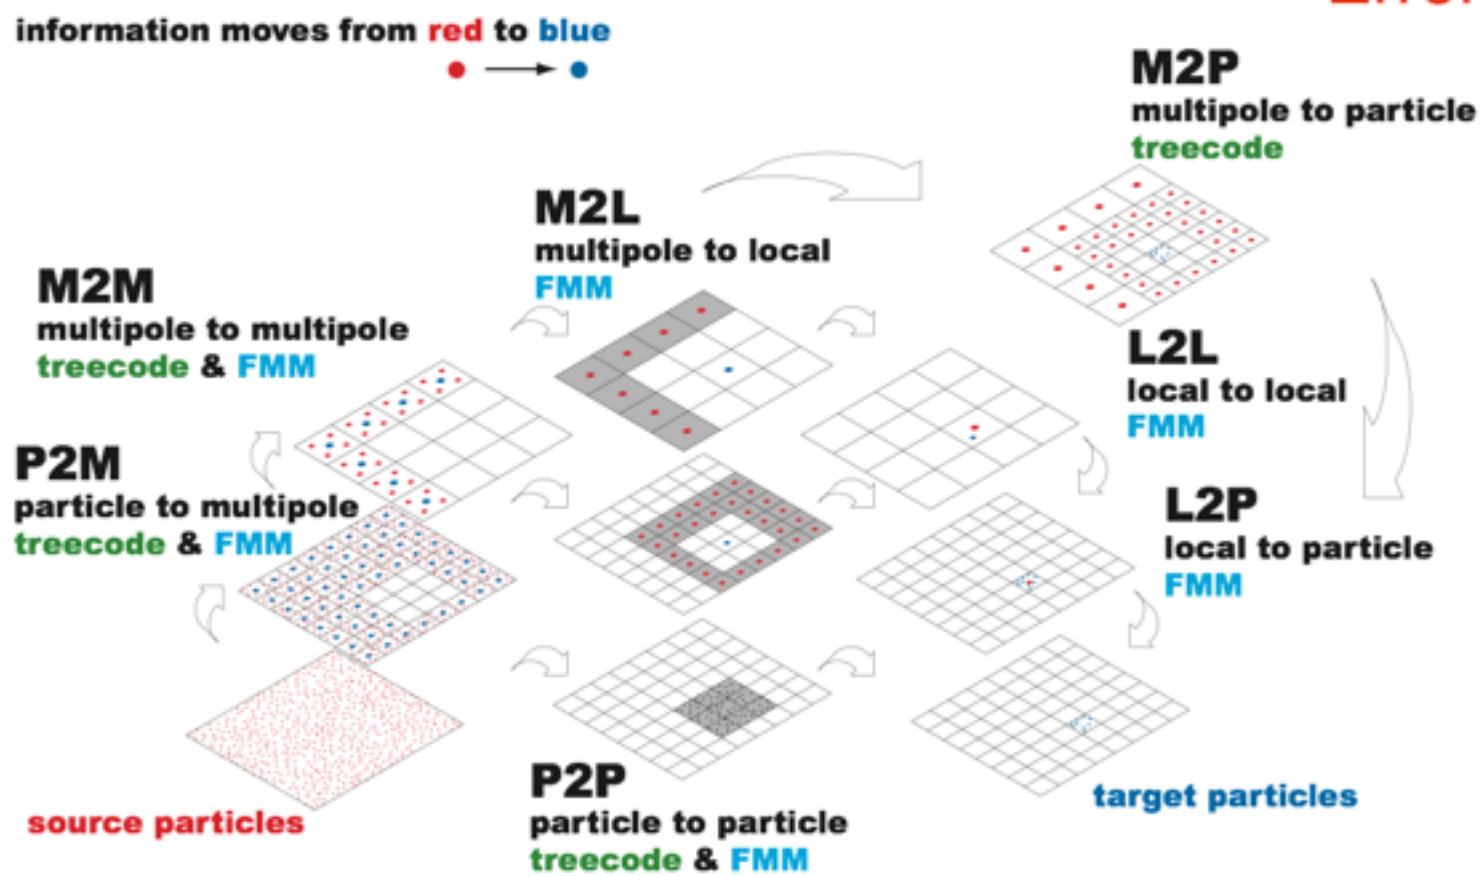


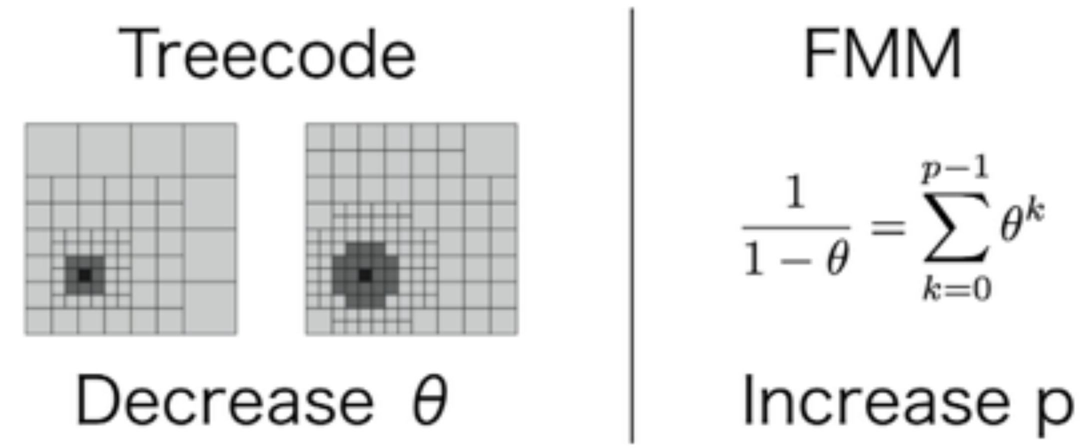


$$\frac{1}{1-\theta} = \sum_{k=0}^{p-1} \theta^k$$

| | $p = 3$ | $p = 4$ | $p = 5$ | $p = 6$ |
|---|---|---|---|---|
| $Err = 10^{-2}$ | $\theta = 1.00$ $time = 0.016s$ | $\theta = 1.18$ $time = 0.012s$ | $\theta = 1.23$ $time = 0.015s$ | $\theta = 1.24$ $time = 0.026s$ |
| $Err = 10^{-3}$ | $\theta = 0.67$ $time = 0.036s$ | $\theta = 0.78$ $time = 0.027s$ | $\theta = 0.91$ $time = 0.024s$ | $\theta = 0.94$ $time = 0.038s$ |
| $Err = 10^{-4}$ | $\theta = 0.30$ $time = 0.22s$ | $\theta = 0.49$ $time = 0.085s$ | $\theta = 0.62$ $time = 0.071s$ | $\theta = 0.70$ $time = 0.073s$ |
| $Err = 10^{-5}$ | $\theta = 0.12$ $time = 1.38s$ | $\theta = 0.20$ $time = 0.59s$ | $\theta = 0.36$ $time = 0.21s$ | $\theta = 0.45$ $time = 0.21s$ |

# Spatially Varying "h" or "p"

Corrected Article: "An error-controlled fast multipole method" [J. Chem. Phys. 131, 244102 (2009)]

Holger Dachsel[a]
*Institute for Advanced Simulation, Jülich Supercomputing Centre, Forschungszentrum Jülich, 52425 Jülich, Germany*

Spatially varying rank

Journal of Computational Physics **179,** 27–42 (2002)
doi:10.1006/jcph.2002.7026

A Hierarchical $\mathcal{O}(N)$ Force Calculation Algorithm

Walter Dehnen

*Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*
E-mail: dehnen@mpia.de

Spatially varying admissibility

↓

SIMD friendly

# Conclusions

| Reference | Processes | | Data per Process | | Communication complexity |
|---|---|---|---|---|---|
| Teng | $\mathcal{O}(P)$ | | $\mathcal{O}\left((N/P)^{2/3}(\log N + \mu)^{1/3}\right)$ | | $\mathcal{O}\left(P(N/P)^{2/3}(\log N + \mu)^{1/3}\right)$ |
| Lashuk *et al.* | $\mathcal{O}(\sqrt{P})$ | | $\mathcal{O}\left((N/P)^{2/3}\right)$ | | $\mathcal{O}\left(\sqrt{P}(N/P)^{2/3}\right)$ |
| Ibeid *et al.* | Global | Local | Global | Local | Global + Local |
| | $\mathcal{O}(\log P)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ | $\mathcal{O}\left((N/P)^{2/3}\right)$ | $\mathcal{O}\left(\log P + (N/P)^{2/3}\right)$ |

- We have proved a new upper bound for the communication complexity of FMM
- The dual tree traversal allows hierarchical methods to use adaptive admissibility conditions, which results in fine grain uniformity that is easier to vectorize

# Title