A Comparison of two *N*-Body Algorithms

Girija Narlikar and Guy Blelloch

School of Computer Science

Carnegie Mellon University

The Third DIMACS International Algorithm
Implementation Challenge
October 17-18, 1994

Outline of this talk

- About the algorithms
- Motivation for this work
- Estimation of operation counts
- Comparison of operation counts
- Conclusions

The Barnes-Hut algorithm

1. Tree construction

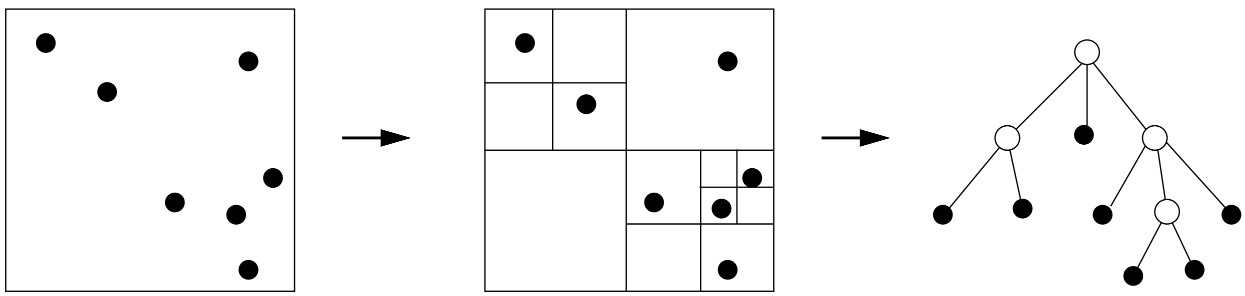

2. Force Calculation

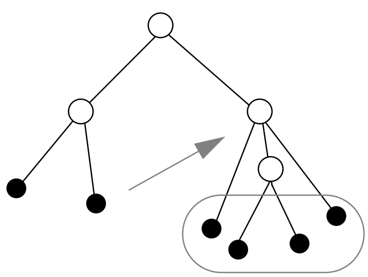

well-separatedness :

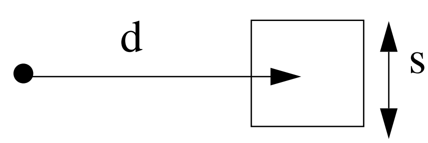

$$\frac{s}{d} < \theta$$

Variable parameter : θ

O(NlogN) for uniform distributions

Greengard's Fast Multipole Method (FMM)

Three phases :

- Bottom up : forming multipole expansions
- Top down : forming local expansions (cell - cell)
- Direct interactions with neighbours (particle - particle)

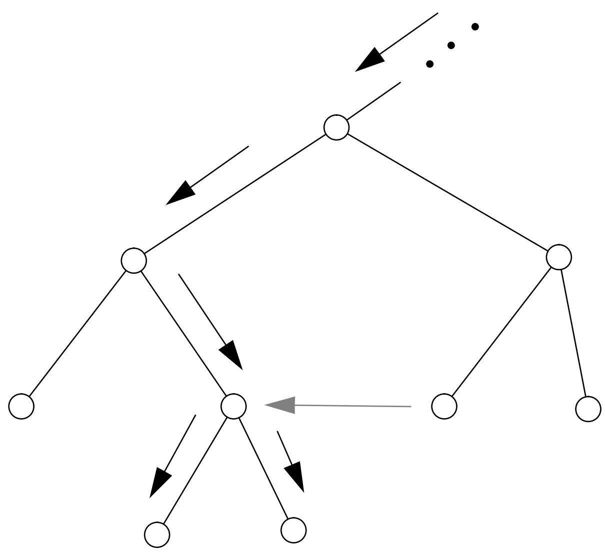

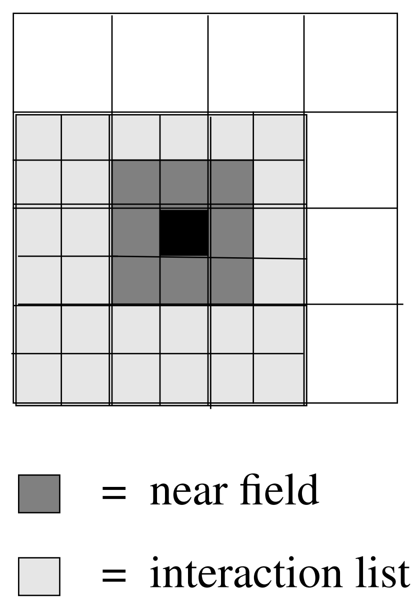

Variable parameter : p

O(N) for uniform distributions

Comparing the two algorithms

Why use floating point operation counts?

- Run-times are machine dependent
- Communication structures are similar
- Both scale well for parallel implementations

	Non-Computational Overheads
Warren-Salmon	15%
Bhatt-Liu	9 %
Zhao	3.4 - 21.5 %

Outline of Experiments

- Dependence of accuracy on θ and *p*
- Expression for #interactions as a function of θ and N for Barnes-Hut
- Comparison of total operation counts of FMM and Barnes Hut for larger N

Experimental Errors

Barnes-Hut :

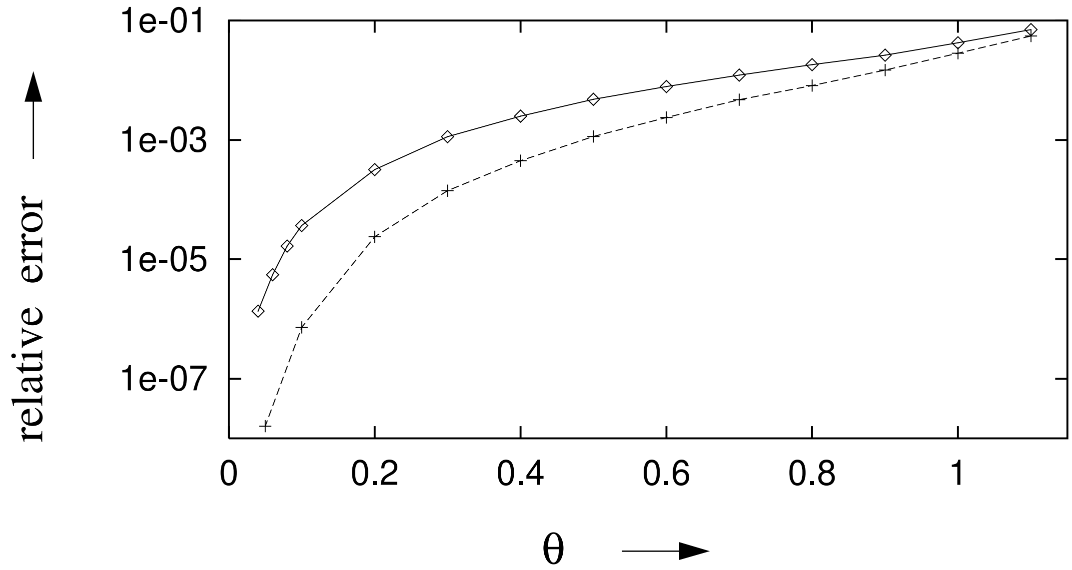

FMM :

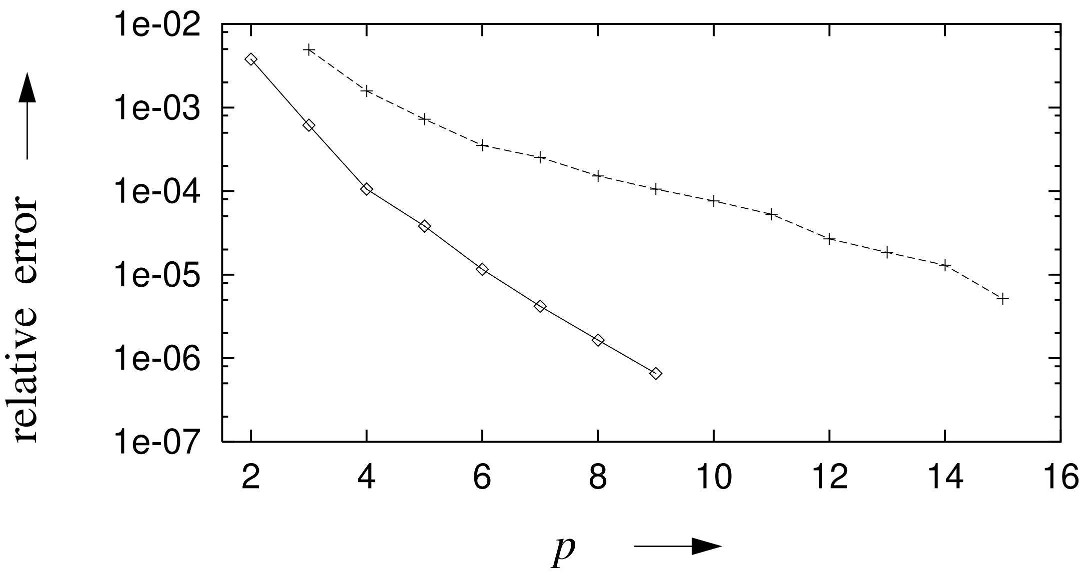

Operation Count for the Barnes-Hut Algorithm

For a random distribution : **f (θ) N logN + ?**

= b(θ) N logN - a(θ) N

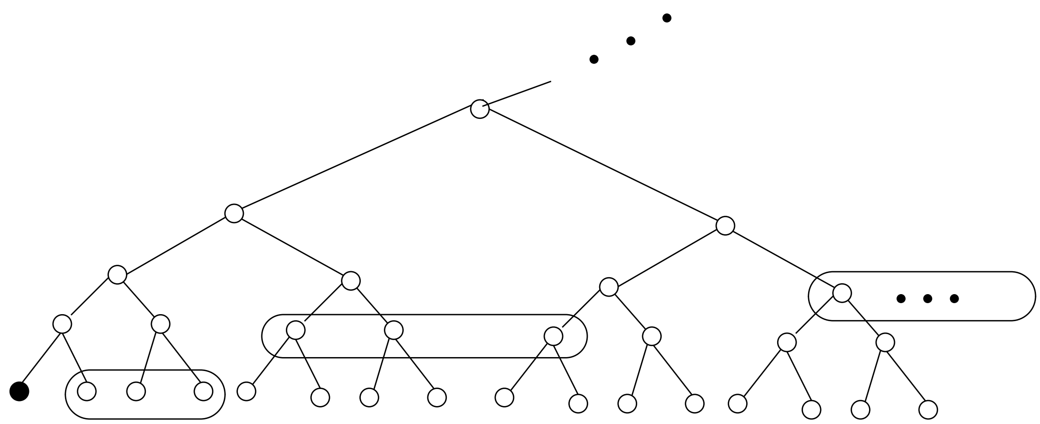

Cell of size **s** at distance **d** interacts directly with particle P if **d > s/θ**

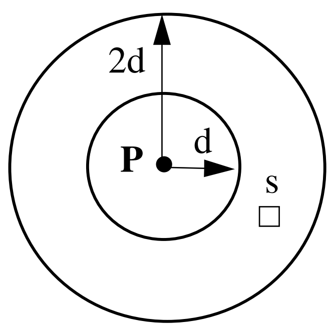

Number of cells of size **s** interacting with P =

$$\mathbf{f(\theta)} = \frac{\text{Volume enclosing these cells}}{\text{Volume of each cell}} = \frac{4/3\pi[(2s/\theta)^3 - (s/\theta)^3]}{s^3} = \mathbf{28\pi/3\theta^3}$$

Levels in the tree = 0,1,2,3.... (starting from the leaves)

particles in a cell at level $\boldsymbol{i}$ $= \mathbf{8}^{i}$

$$\mathbf{N} = \sum_{i=0}^{i<=L} \mathbf{f}(\theta) \cdot \mathbf{8}^{i} \quad \Rightarrow \quad \mathbf{L} \approx \mathbf{log}_{8}\, [\mathbf{N} / \mathbf{f}(\theta)]$$

$\therefore$ Total number of interactions = $\mathbf{N} \cdot \mathbf{f}(\theta) \cdot \mathbf{log}_{8}\, [\mathbf{N} / \mathbf{f}(\theta)]$

$$= \underbrace{\mathbf{f}(\theta)}_{\mathbf{b}(\theta)}\, \mathbf{N}\, \mathbf{log}_{8}(\mathbf{N}) - \underbrace{\mathbf{f}(\theta)\, \mathbf{log}_{8}(\mathbf{f}(\theta))}_{\mathbf{a}(\theta)}\, \mathbf{N}$$

Measured and Derived Values of b(θ)

θ	f(θ)	b(θ)	a(θ)/b(θ)
1.0	29.28	26.8	1.526
0.7	85.44	78	2.168
0.6	135.3	120.6	2.29
0.5	234.2	209.8	2.63
0.3	1082	840	3.86
0.2	3737	3077.5	3.66

Interactions vs. N

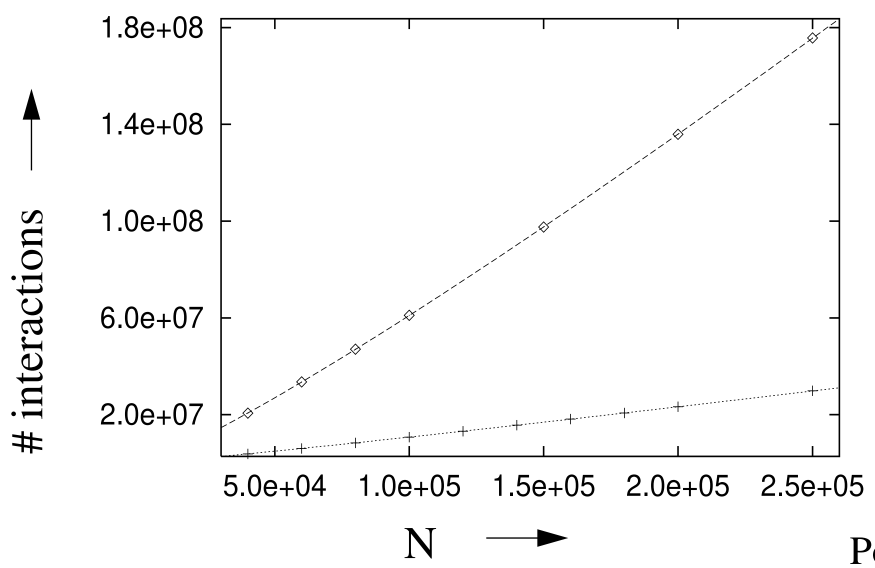

Points : Measured values

Lines : fitted curves

Comparison of Operation Counts for fixed errors

At High Accuracy : (RMS error < 2×10^{-5})

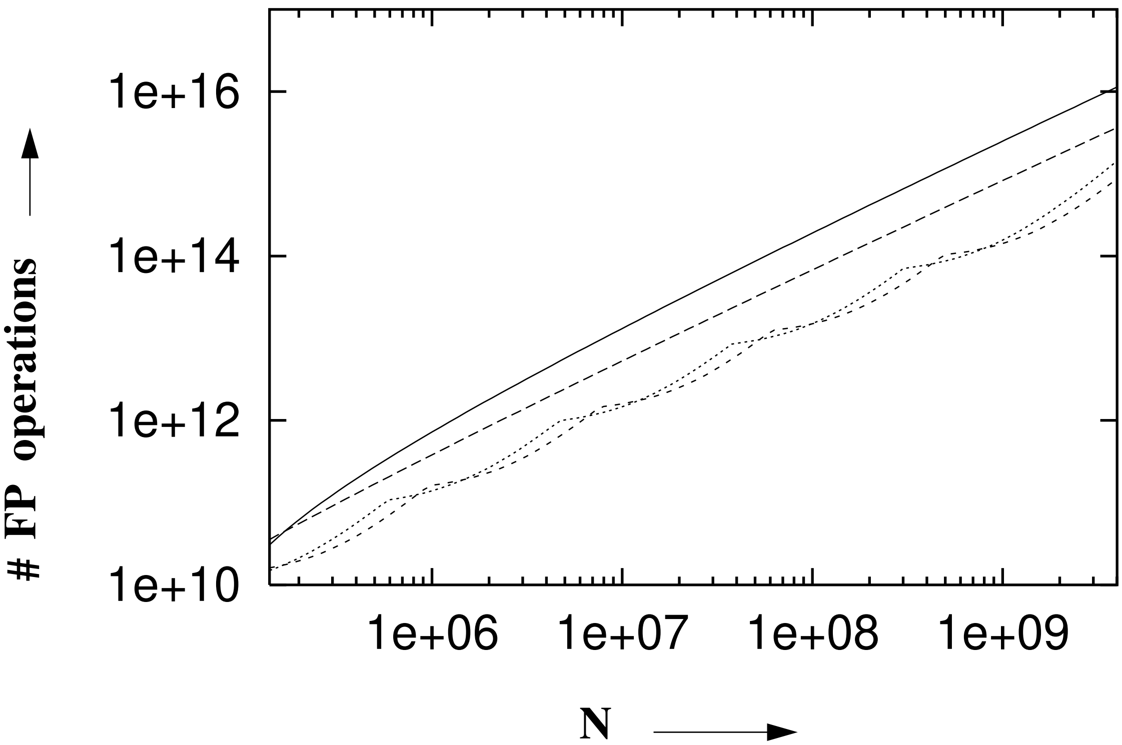

Ratio of Operation Counts to FMM_1

(at high accuracy)

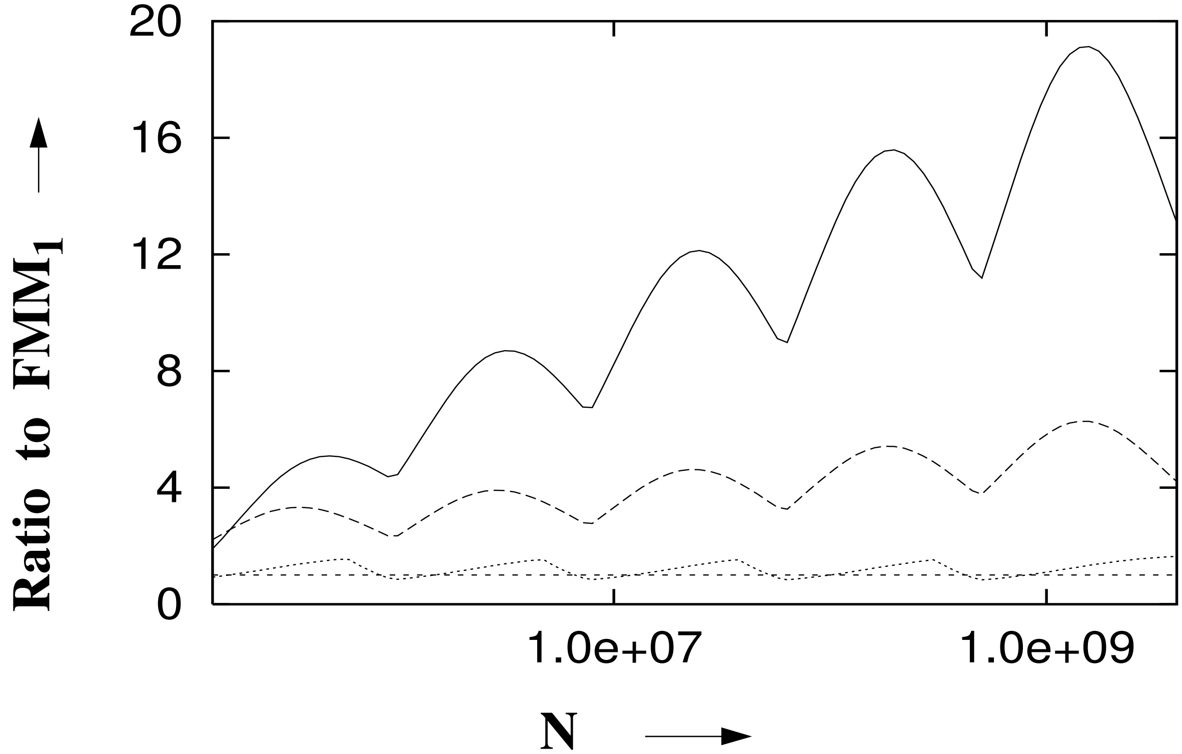

Ratio of operation counts to FMM_1

At Low Accuracy : (RMS error < 10^{-3})

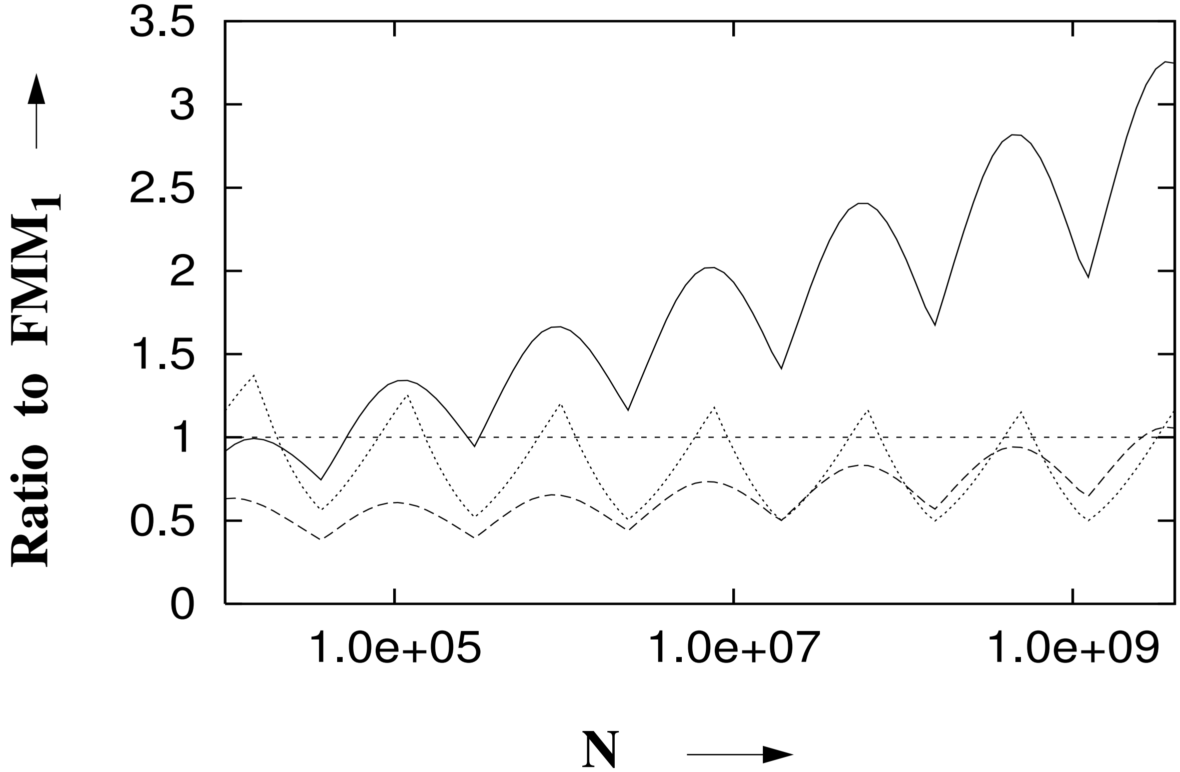

Conclusions

- Linear term in Barnes-Hut cannot be ignored
- The FMM is more complicated, has bigger hidden constants
- Both do comparable amounts of work for uniform distributions

Current / Future Work

- Performance for different data distributions
- Effects of various optimizations

Work on Plummer Model