


Parallel Planar Closest Point Pair

Plan-of-action and Results



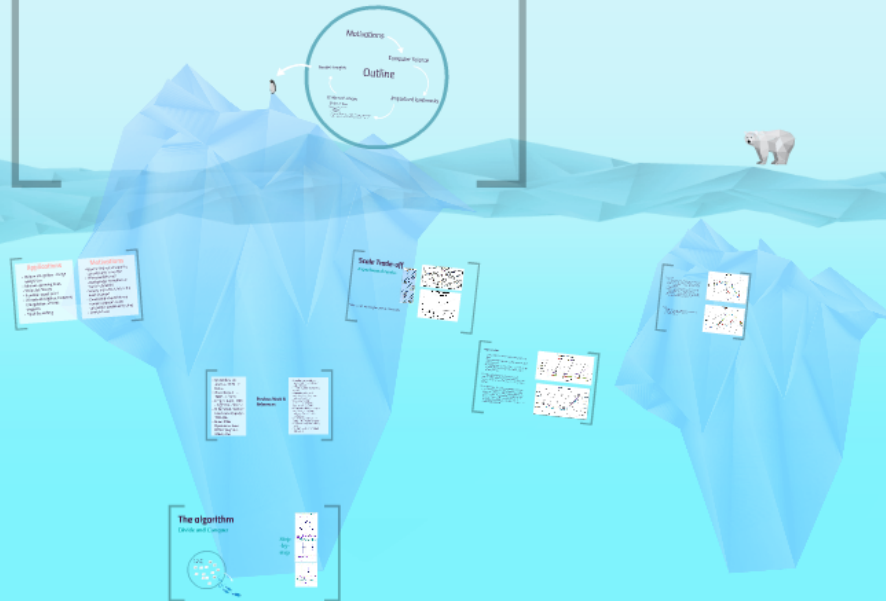
P - 4 | CSE633 Spring 2014

Parallel Planar Closest Point Pair

Angad Gadre & Piyush Sankla
Email: angadgad@buffalo.edu

Parallel Planar Closest Point Pair

Plan-of-action and Results



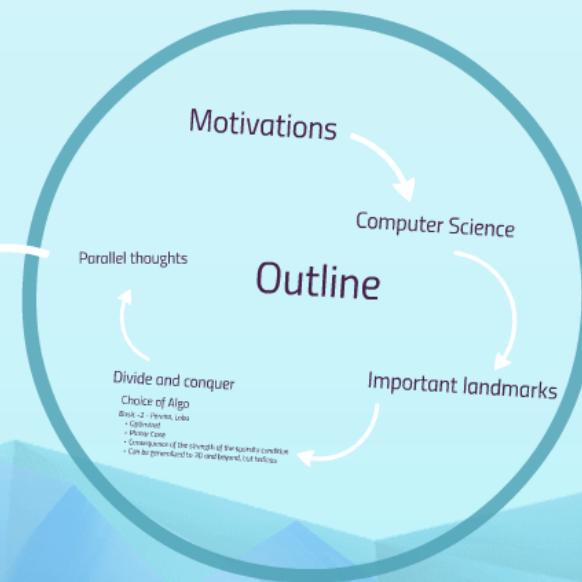
P - 4 | CSE633 Spring 2014

Parallel Planar Closest Point Pair

Angad Gadre & Piyush Sankla
Email: angadgad@buffalo.edu

Parallel Planar Closest Point Pair

Plan-of-action and Results



Motivations



Computer Science



Important landmarks

Outline



Divide and conquer

Choice of Algo

Basic -2 - Pereira, Lobo

- Optimized
- Planar Case
- Consequence of the strength of the sparsity condition
- Can be generalized to 3D and beyond, but tedious



Parallel thoughts



Applications

- Pattern recognition - Image recognition
- Minimal spanning trees
- Molecular biology
- Iterative closest point
- All nearest neighbor, Delaunay triangulation, Voronoi diagrams
- Tip of the iceberg

Motivations

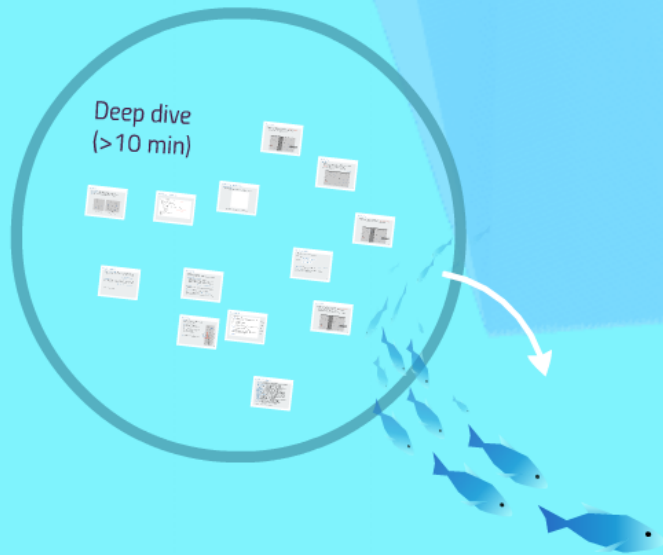
- Want to find out who are the closest people in my life?
- Who resembles me?
Doppelganger or multiverse cousin? (Dynamic)
- How to reach Davis hall in the least distance?
- Understand why divide and conquer approaches are suitable for parallel computing
- Amdahl's law

- Jon Bentley, Ian Shamos, 1976 - 7 Points
- Zhou, Xiong, Zhu, 1998 - 4 points
- Jiang, Gillespie, 2007 - log factor - Basic 2
- Miller, Stout, Mesh of Processors $O(\sqrt{n})$, 1984, 86
- Boxer, Miller - Dynamic versions, PRAM $O(\log^2 n)$. Widely cited

**Previous
Referen**


The algorithm

Divide and Conquer



Step-by-step


- We partition S into S_1, S_2 by vertical line ℓ defined by median x -coordinate in S .
- Recursively compute closest pair distances δ_1 and δ_2 . Set $\delta = \min(\delta_1, \delta_2)$.
- Now compute the closest pair with one point each in S_1 and S_2 .



- In each candidate pair (p, q) , where $p \in S_1$ and $q \in S_2$, the points p, q must both lie within δ of ℓ .
- Consider a point $p \in S_1$. All points of S_2 within distance δ of p must lie in a $\delta \times 2\delta$ rectangle R .




- How many points can be inside R if each pair is at least δ apart?
- In 2D, this number is at most 6!
- So, we only need to perform $6 \times n/2$ distance comparisons!
- In order to determine at most 6 potential mates of p , project p and all points of P_2 onto line ℓ .




- Pick out points whose projection is within δ of p ; at most six.
- We can do this for all p , by walking sorted lists of P_1 and P_2 , in total $O(n)$ time.
- The sorted lists for P_1, P_2 can be obtained from pre-sorting of S_1, S_2 .
- Final recurrence is $T(n) = 2T(n/2) + O(n)$, which solves to $T(n) = O(n \log n)$.

Step -by- step

- We partition S into S_1, S_2 by vertical line ℓ defined by median x -coordinate in S .
- Recursively compute closest pair distances δ_1 and δ_2 . Set $\delta = \min(\delta_1, \delta_2)$.
- Now compute the closest pair with one point each in S_1 and S_2 .




- In each candidate pair (p, q) , where $p \in S_1$ and $q \in S_2$, the points p, q must both lie within δ of ℓ .
- Consider a point $p \in S_1$. All points of S_2 within distance δ of p must lie in a $\delta \times 2\delta$ rectangle R .



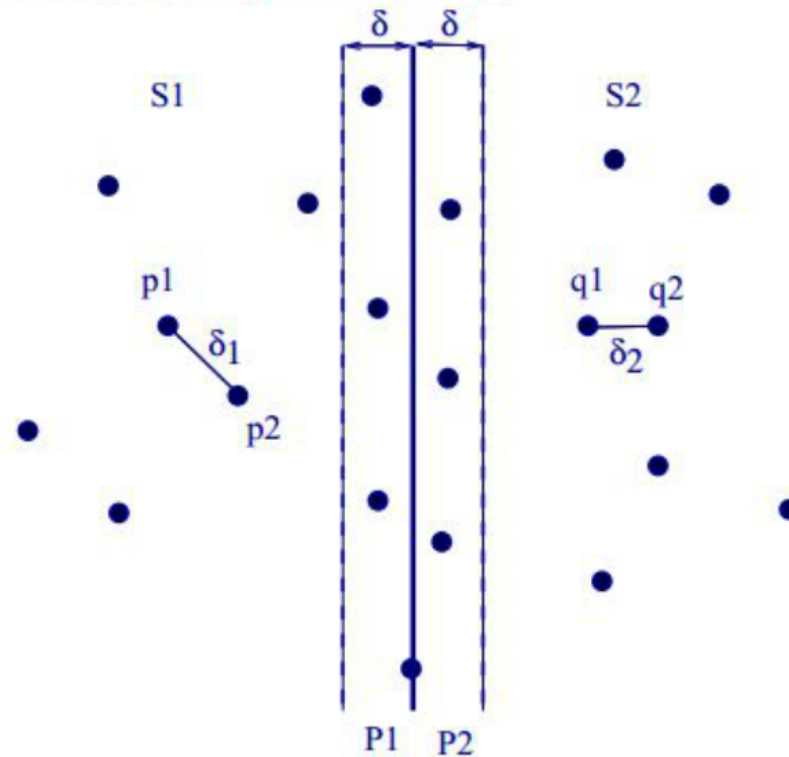
- How many points can be inside R if each pair is at least δ apart?
- In 2D, this number is at most 6!
- So, we only need to perform $6 \times n/2$ distance comparisons!

- In order to determine at most 6 potential mates of p , project p and all points of P_2 onto line ℓ .




- Pick out points whose projection is within δ of p ; at most six.
- We can do this for all p , by walking sorted lists of P_1 and P_2 , in total $O(n)$ time.
- The sorted lists for P_1, P_2 can be obtained from pre-sorting of S_1, S_2 .
- Final recurrence is $T(n) = 2T(n/2) + O(n)$, which solves to $T(n) = O(n \log n)$.

- We partition S into S_1, S_2 by vertical line ℓ defined by median x -coordinate in S .
- Recursively compute closest pair distances δ_1 and δ_2 . Set $\delta = \min(\delta_1, \delta_2)$.
- Now compute the closest pair with one point each in S_1 and S_2 .




- In each candidate pair (p, q) , where $p \in S_1$ and $q \in S_2$, the points p, q must both lie within δ of ℓ .

- Consider a point $p \in S_1$. All points of S_2 within distance δ of p must lie in a $\delta \times 2\delta$ rectangle R .





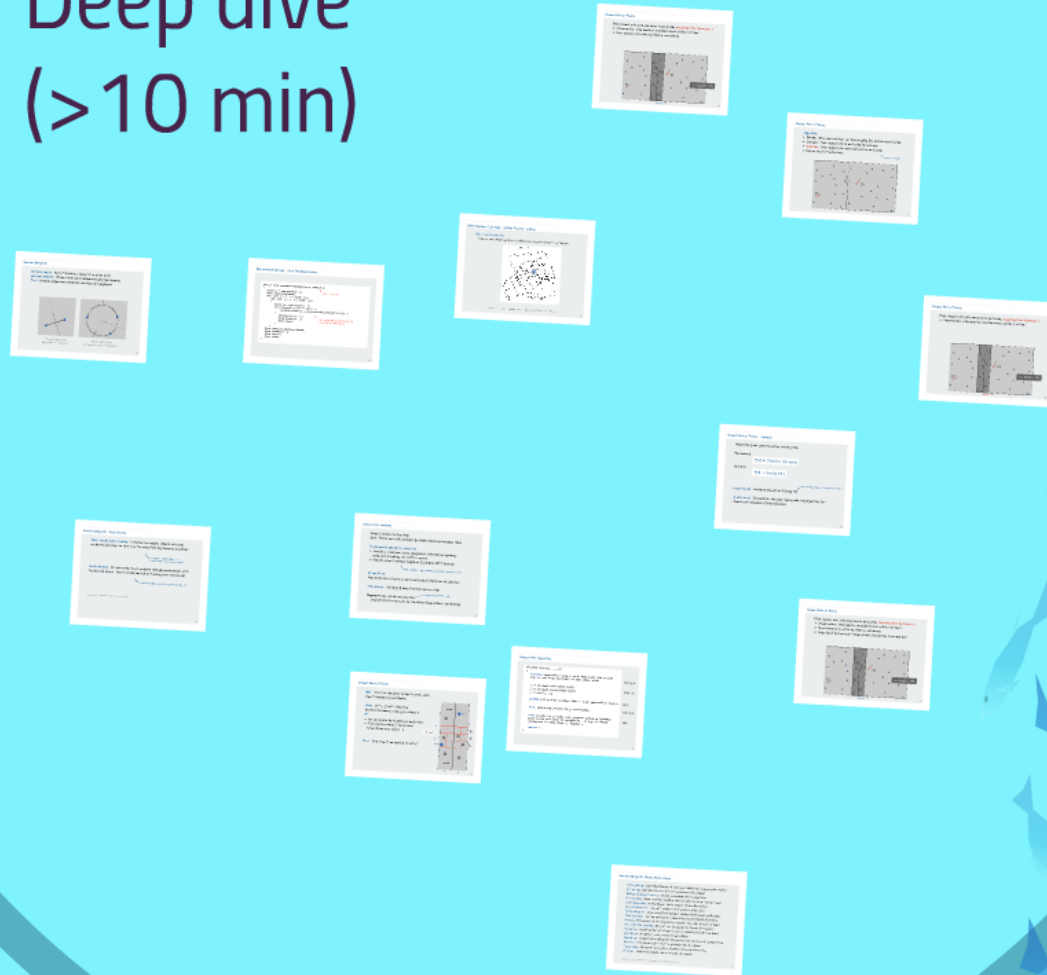
- How many points can be inside R if each pair is at least δ apart?
- In 2D, this number is at most 6!
- So, we only need to perform $6 \times n/2$ distance comparisons!

- In order to determine at most 6 potential mates of p , project p and all points of P_2 onto line ℓ .



- Pick out points whose projection is within δ of p ; at most six.
- We can do this for all p , by walking sorted lists of P_1 and P_2 , in total $O(n)$ time.
- The sorted lists for P_1, P_2 can be obtained from pre-sorting of S_1, S_2 .
- Final recurrence is $T(n) = 2T(n/2) + O(n)$, which solves to $T(n) = O(n \log n)$.

Deep dive (> 10 min)



Brute Force

Points	Time (sec)
100	0.000018
1000	0.001757
10000	0.164719
100000	16.323205
1000000	Too long

Divide and Conquer


Points	Time (sec)
100	0.000095
1000	0.000559
10000	0.006099
100000	0.081928
1000000	1.110317
1000000	2.550909
1000000	5.810905
1000000	7.703420
1000000	12.714972
1000000	14.095037
1000000	16.361776
1000000	19.348326

Timing Data for "Closest"

	Number of Points				
Machines	1000	5000	10000	50000	250000
1	.82	3.63	8.15	54.83	470.87
3	.90	2.90	5.61	33.35	259.15
7	1.17	2.49	4.60	23.00	166.65

SPEEDUP CURVES

NO LOAD BALANCING



Scale Trade-off

Experimental results

Brute Force	
#Points	Time (sec)
100	0.000018
1000	0.001757
10000	0.164719
100000	16.323205
1000000	Too long

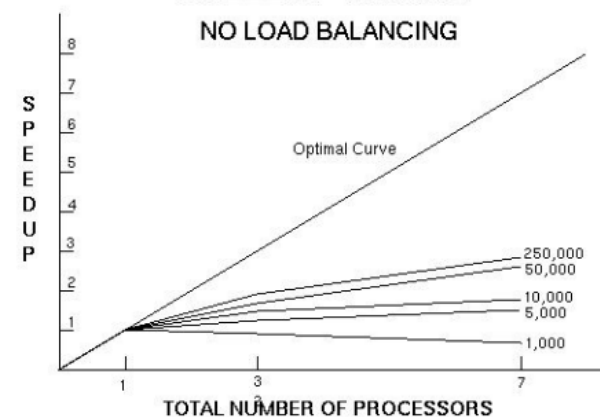
Divide and Conquer	
#Points	Time (sec)
100	0.000095
1000	0.000559
10000	0.006099
100000	0.081928
1000000	1.110317
2000000	2.550909
4000000	5.810905
5000000	7.703420
6000000	12.714972
8000000	14.095037
9000000	16.361776
10000000	19.348326

Timing Data for "Closest"

Machines	Number of Points				
	1000	5000	10000	50000	250000
1	.82	3.63	8.15	54.83	470.87
3	.90	2.90	5.61	33.35	259.15
7	1.17	2.49	4.60	23.00	166.65

SPEEDUP CURVES

NO LOAD BALANCING



*John Thrall, Washington and Lee University

Approach for MPI

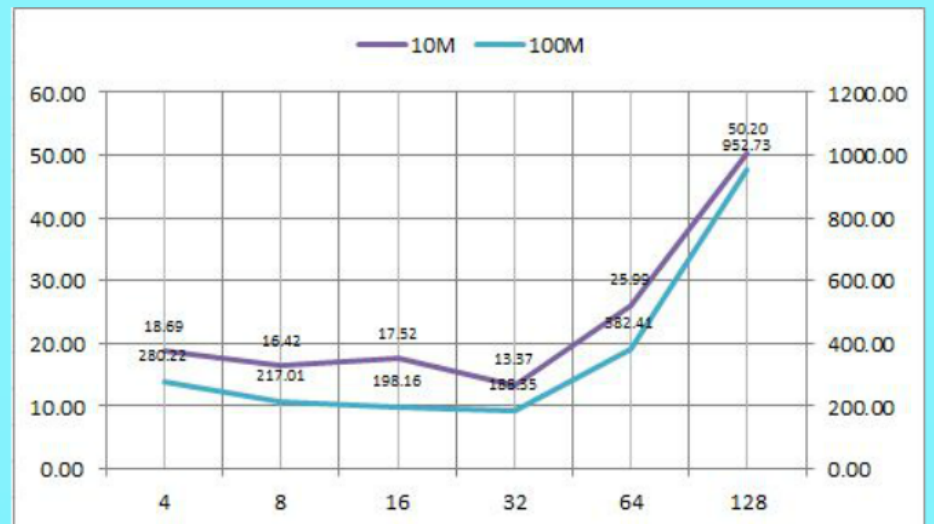
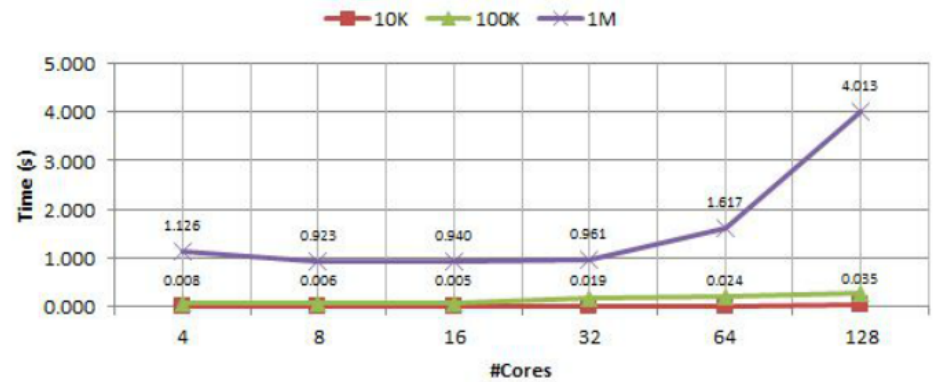
- Each process is given a section of the sequence of data.
- This is done by performing a scatter operation by the master process.
- Each process works on its subsequence individually using the P4 algorithm explained and sends its output for merge.

The major advantage of this technique is that communication time is kept low by sending only a sub-sequence to each process. But this is not an optimized solution.

The disadvantages of such a technique are as follows:

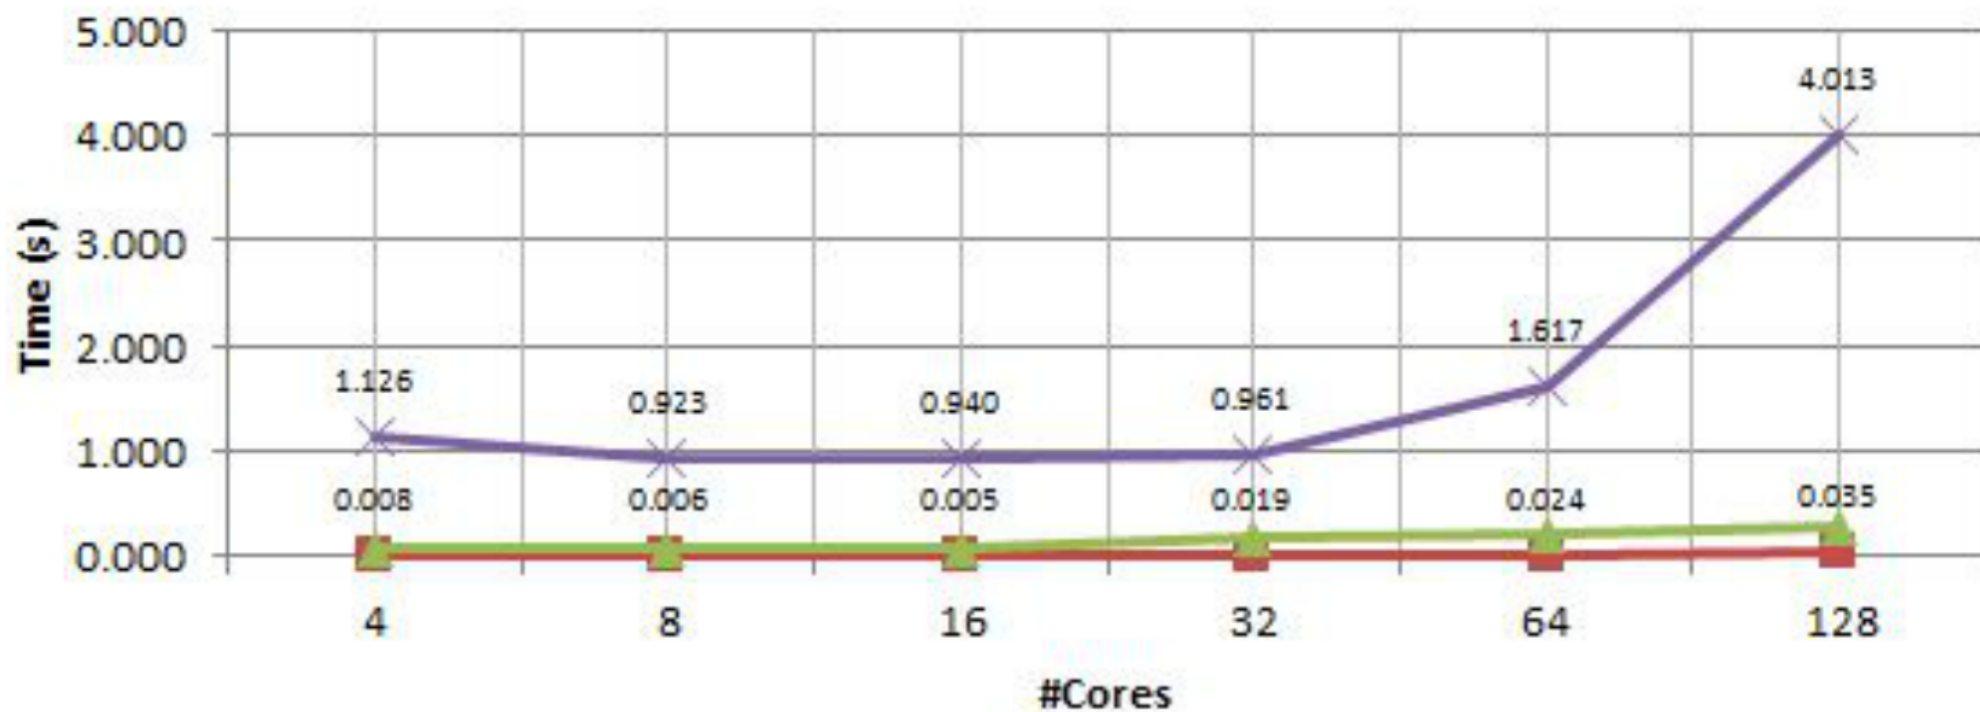
- Load balancing is not achieved since a particular process may finish its process and wait for merge operation while other processes are still busy. This leads to idle process time and is not efficient.
- Merge operations performed for each sub-sequence is also computationally.

Running Time vs PE




Running Time vs PE


■ 10K ▲ 100K × 1M



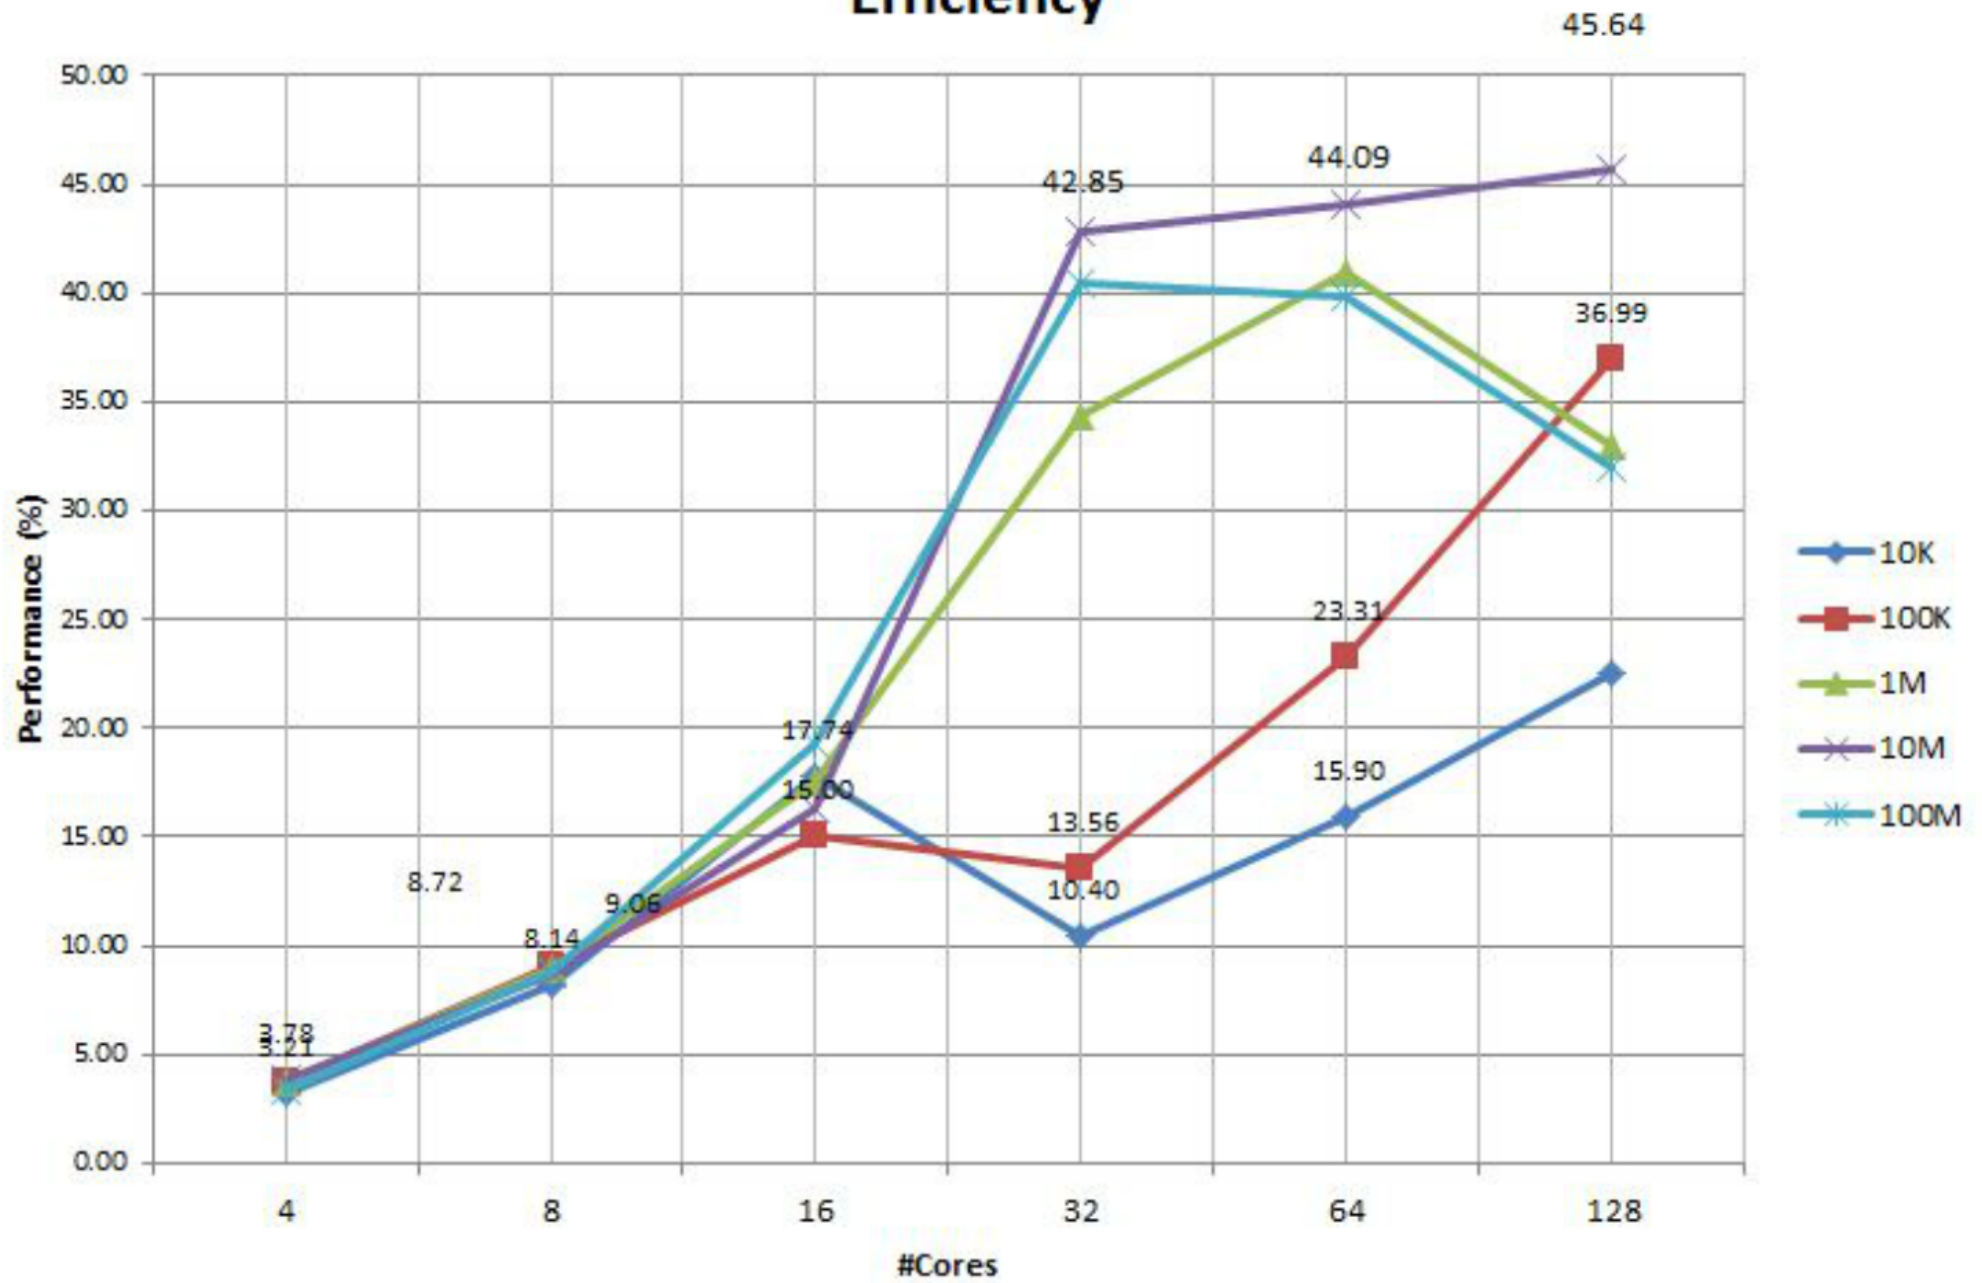
— 10M — 100M



Speedup

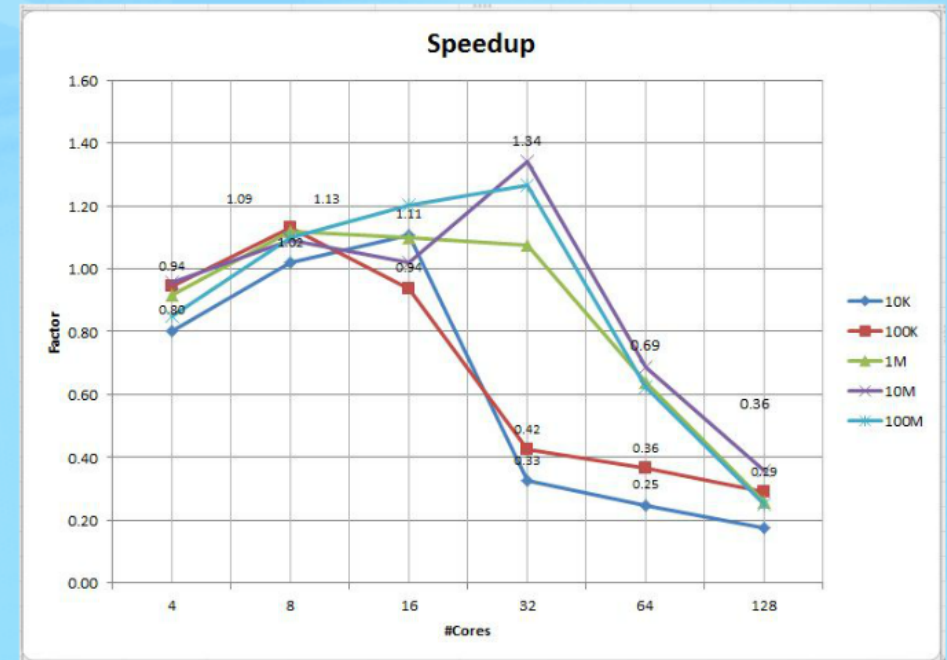


Efficiency



Conclusions

- In engineering contexts, efficiency is more often used for graphs than speedup, since all of the area in the graph is useful (whereas in a speedup curve 1/2 of the space is wasted)
- it is easy to see how well parallelization is working
- there is no need to plot a "perfect speedup" line
- In marketing contexts, speedup curves are more often used, largely because they go up and to the right and thus appear better to the less-informed.



Future

- Complete divide and conquer using hyperquicksort
- Implement pseudo EM like K-means
- Use Slurm effectively



s Work & nces

- Algorithms: Sequential and parallel, 2nd Ed; Russ Miller & Laurence Boxer
- Parallel algorithms in geometry, Goodrich
- Sequential and parallel algorithms for k closest pair problem; Lenhof, Smid
- A survey of Parallel computational geometry algorithms; Hehne, Sack
- An optimized divide and conquer algorithm for closest pair problem in the planar case; Pereira, Lobo
- Geometric Algorithms; Robert Sedgwick, Princeton University
- Closest Pair (PPT); Subhash Suri, UCSB
- Parallel Computing 101; Stout, Jablonowski