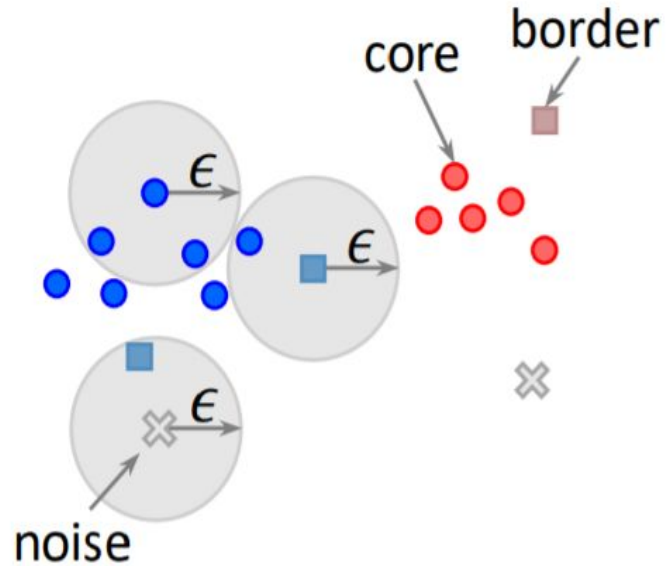


Theoretically-Efficient and Practical Parallel DBSCAN

Yiqiu Wang, Yan Gu, Julian Shun

DBSCAN (ϵ , minpts)

Minpts = 3



Previous work

- Parallel
 - Xu et al. PDBSCAN
 - Distributed R-Tree, exact answer
 - Patwary et al. PDSDBSCAN
 - Parallel lock based union-find, approximate without guarantee
 - Gotz et al. HPDBSCAN
 - Data partition, process locally and then merge results, exact answer
 - RP-DBSCAN
 - Distributed map-reduce, approximate without guarantee
 - ...
- Serial
 - Gunawan et al.
 - Sequential exact DBSCAN with theoretical guarantees
 - Gan and Tao
 - Sequential approximate DBSCAN with theoretical guarantees

Our contribution


- Practical parallel algorithms for 2D exact DBSCAN, and higher dimensional exact and approximate DBSCAN with work bounds matching the best sequential algorithms, and polylogarithmic depth.
- Highly-optimized implementations.
- A comprehensive experimental evaluation
 - 2-38x self-relative speedup (36 cores)
 - 5-33x speedup over best sequential (36 cores)
 - 14-6102x faster than existing parallel implementations
 - Processes largest dataset known for DBSCAN (> 4 billion points) on just one machine

DBSCAN High-level Algorithm (2D)

- Construct cells
- Mark core points
- Cluster core points (cell graph)
- Cluster border points

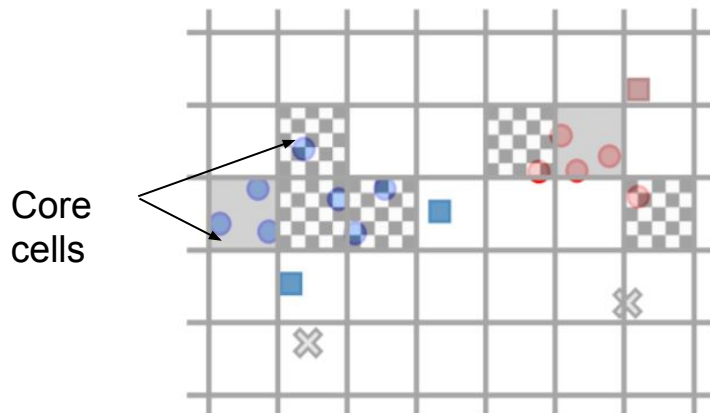
DBSCAN High-level Algorithm (2D)

- **Construct cells**
 - Mark core points
 - Cluster core points (cell graph)
 - Cluster border points
- Side length $\epsilon/\sqrt{2}$
 - $O(n)$ work in expectation
 $O(\log n)$ depth w.h.p.




DBSCAN High-level Algorithm (2D)


- Construct cells
 - **Mark core points**
 - Cluster core points (cell graph)
 - Cluster border points
- For each cell
 - If $|\text{cell}| > \text{minpts}(3)$: all are core points
 - Else: perform range count for each point
 - $O(n \cdot \text{minPts})$ work
 $O(\log n)$ depth w.h.p.



DBSCAN High-level Algorithm (2D)

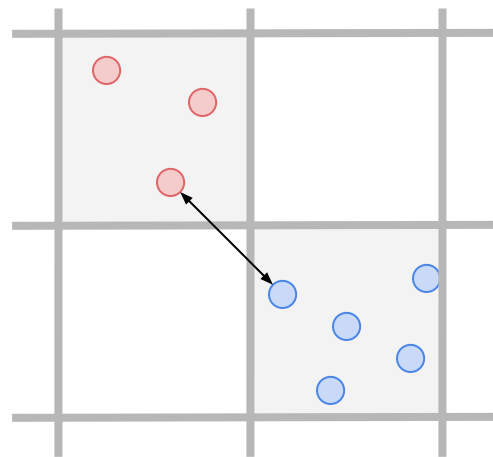
- Construct cells
- Mark core points
- **Cluster core points (cell graph)**
- Cluster border points
- Cell graph + connected components
- Connect “core cells” (cell with ≥ 1 core points) if they contain core points with $< \epsilon$ distance
- **Connected components**
 - $O(n)$ expected work
 - $O(\log n)$ depth w.h.p.





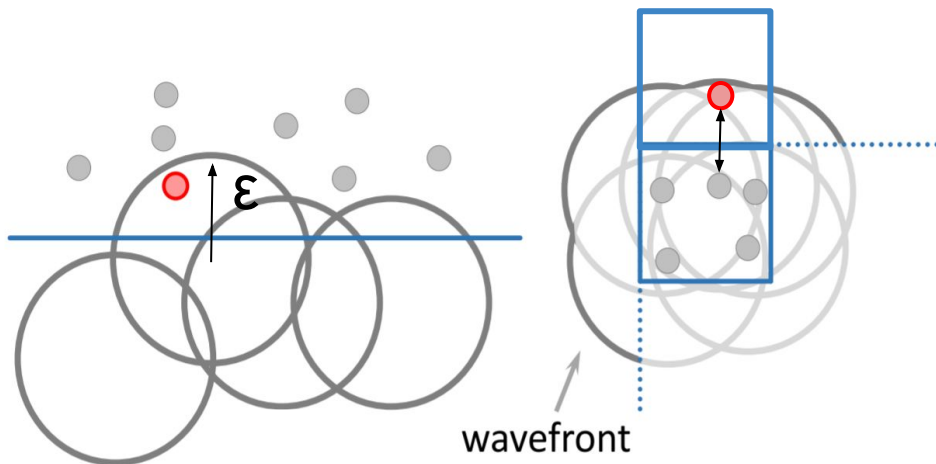
Cell graph - Bichromatic Closest Pair

- Check all pairs of points; pick the minimum
- $O(n^2)$ work
- $O(\log n)$ depth
- Practice
 - Simple optimization to get rid of most points





Cell graph - Unit-spherical emptiness checking (USEC)

- Serial version, Gan and Tao (2015)
- Create “wavefront” between two cells and check for emptiness
- $O(n \log n)$ expected work
- $O(\log^2 n)$ depth w.h.p.
- Practice
 - Sufficient parallelism on grids
 - Serial construction




Cell graph - Delaunay Triangulation

- Serial version, Gan and Tao (2015)
- Connect two cells if they are connected by a Delaunay edge of weight $< \epsilon$
- Reif & Sen (1992)
 - $O(n \log n)$ work
 - $O(\log n)$ depth

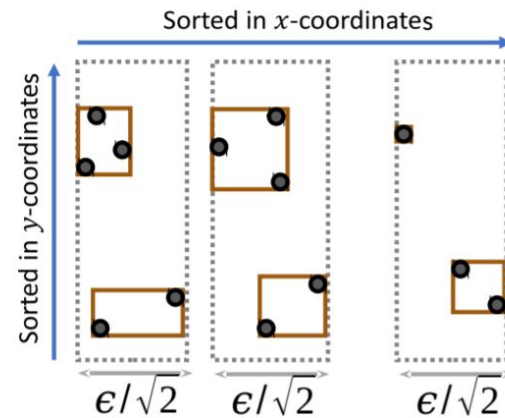


Two types of cells

- Grid cell (Serial version, Gunawan 2013)

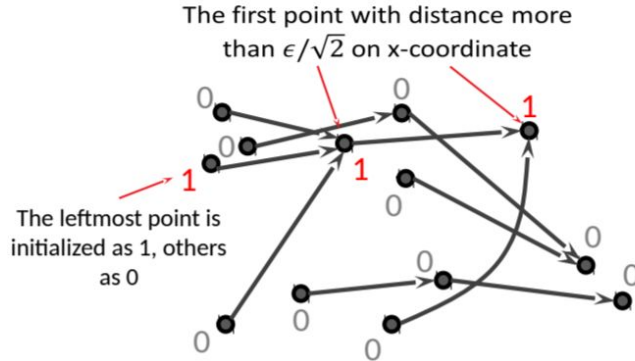



- Box cell (Serial version, Gunawan 2013)



Box cell construction in parallel

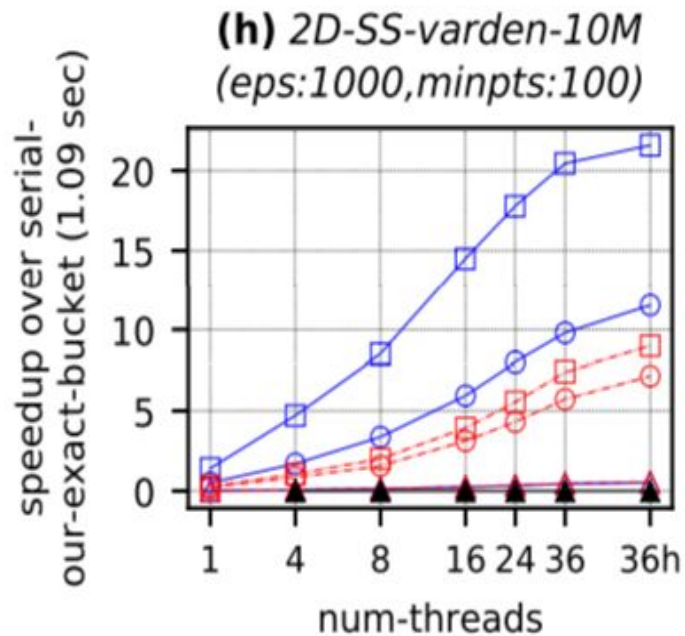
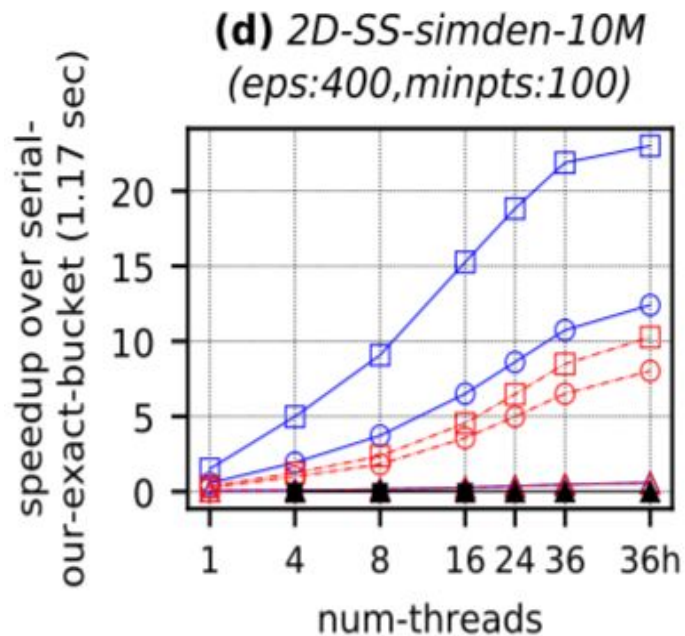
- Vertical columns + Horizontal columns
- Binary search to find next boundary
- Pointer jumping to mark column starts
- $O(\log n)$ depth



Running Times for 2D

Cell construction	Grid based	O(n) expected work; O(log n) depth w.h.p.		
	Box based	O(n log n) work; O(log n) depth		
Mark core		O(n) work; O(log n) depth w.h.p.		
Cell graph		BCP	USEC	Delaunay
		O(n ²) work; O(log n) depth	O(n log n) work; O(log ² n) depth	O(n log n) work; O(log n) depth
Connected components		O(n) expected work; O(log n) depth w.h.p.		
Cluster border		O(n) work; O(log n) depth w.h.p.		
Overall		O(n ²) expected work; O(log n) depth w.h.p.	O(n log n) expected work; O(log ² n) depth w.h.p.	O(n log n) expected work; O(log n) depth w.h.p.

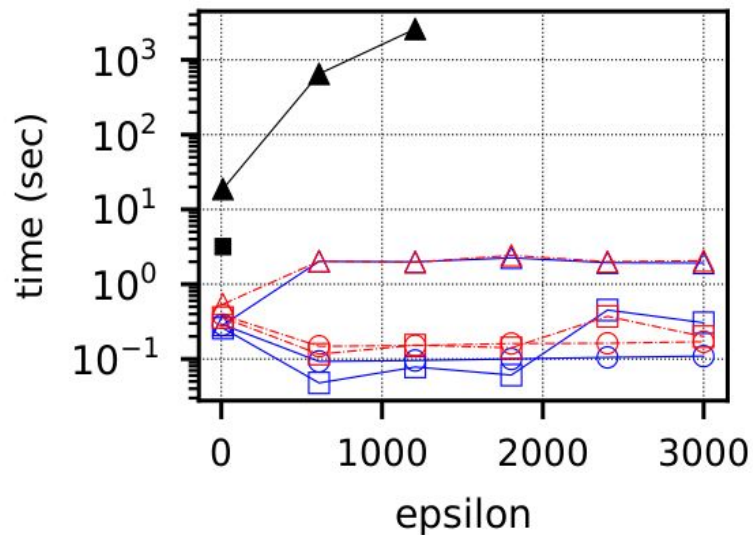
2D Results




2D Results



(a) 2D-SS-simden-10M
(minpts:100)




(e) 2D-SS-vardeen-10M
(minpts:100)




2D Results



(b) 2D-SS-simden-10M
(eps:400)



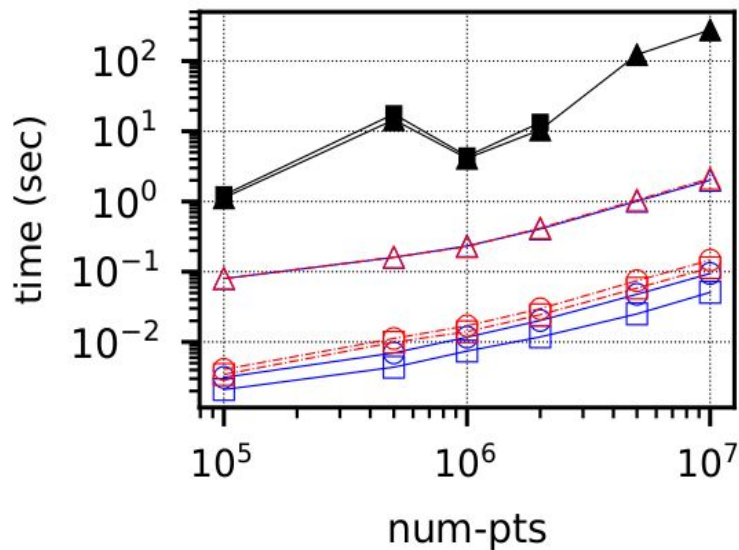
(f) 2D-SS-vardeen-10M
(eps:1000)



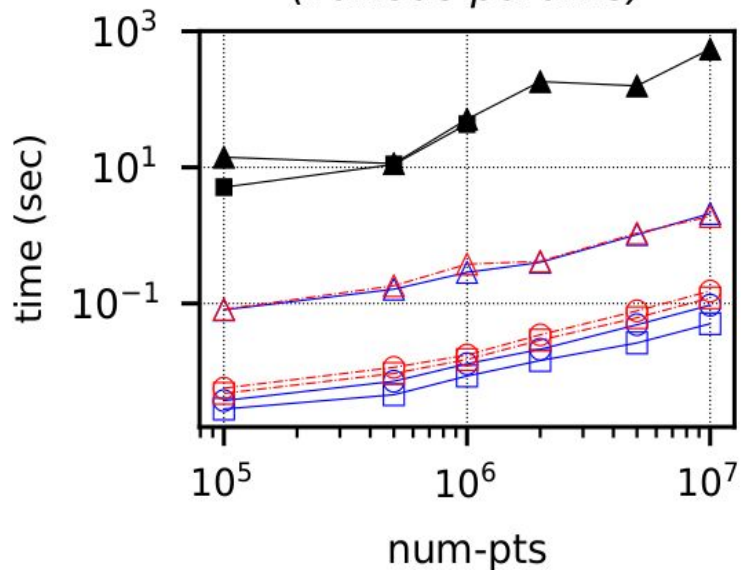
2D Results



(c) 2D-SS-simden-NumPts
(various params)




(g) 2D-SS-varden-NumPts
(various params)



Higher dimensional DBSCAN

- More complicated cell-neighbor finding
- Parallel kd-tree for answer spatial range queries



21 neighbor grids

$\approx 5^d$

Optimization 1 - UnionFind


- Serial version, Gan and Tao (2015)
- BCP & USEC
- Compute CC while constructing the cell graph
- Prunes independent cell-cell connectivity queries

Optimization 2 - Bucketing

- Cells with more points -> more likely connected
 - Connect them early
 - More pruning
- Serial: sort + iterate
- Parallel: sort + bucketing


Cell connectivity using spatial quadtree

- For both exact DBSCAN and approximate DBSCAN




—×— our-approx-qt —+— our-approx — our-best-exact



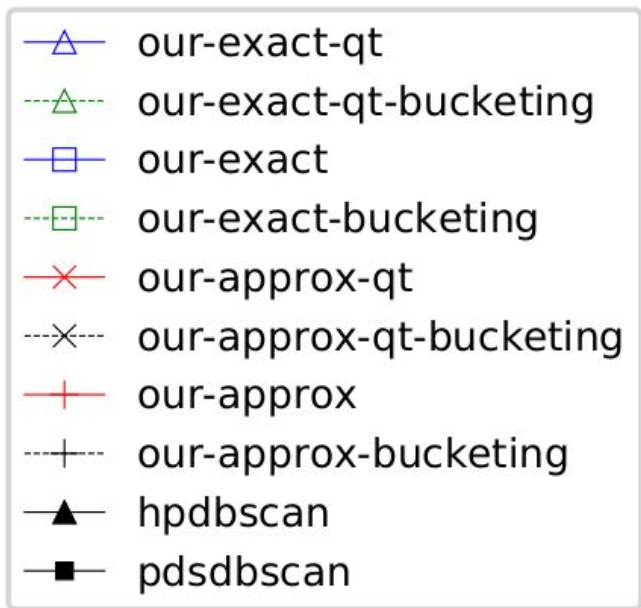
(a) 5D-SS-simden-10M
(minpts:100)



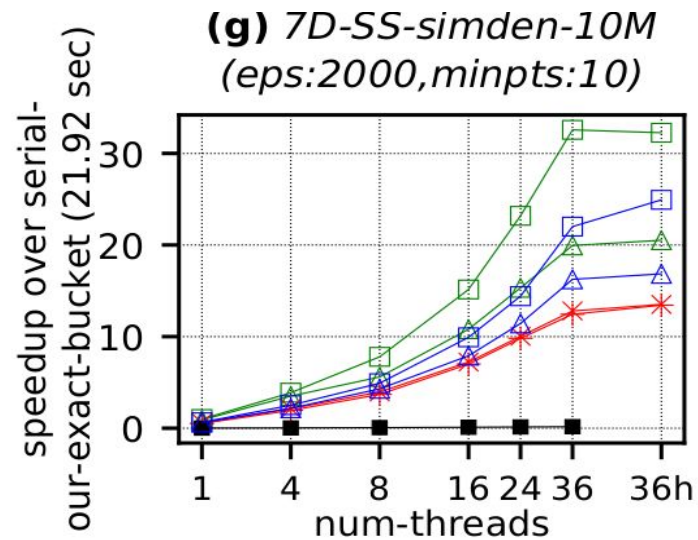

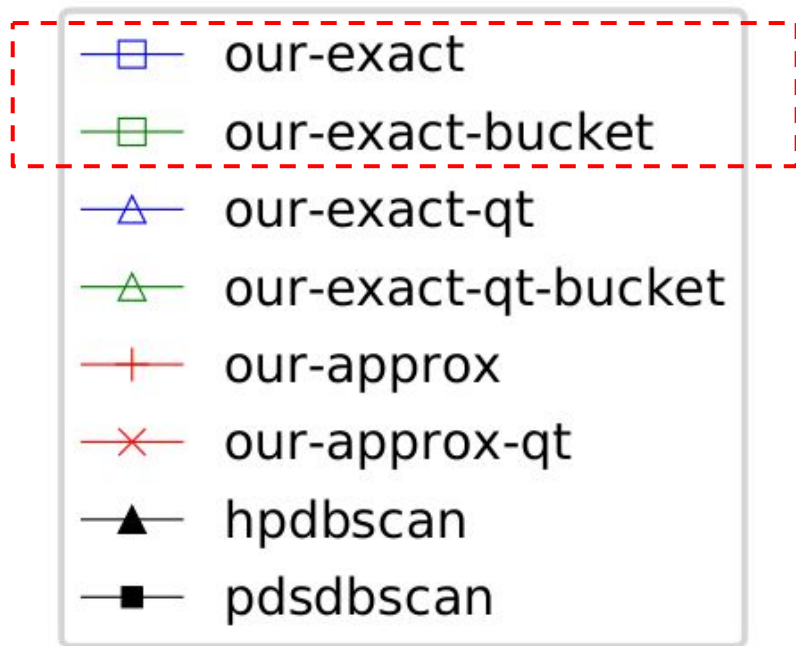
(b) 5D-SS-variden-10M
(minpts:10)



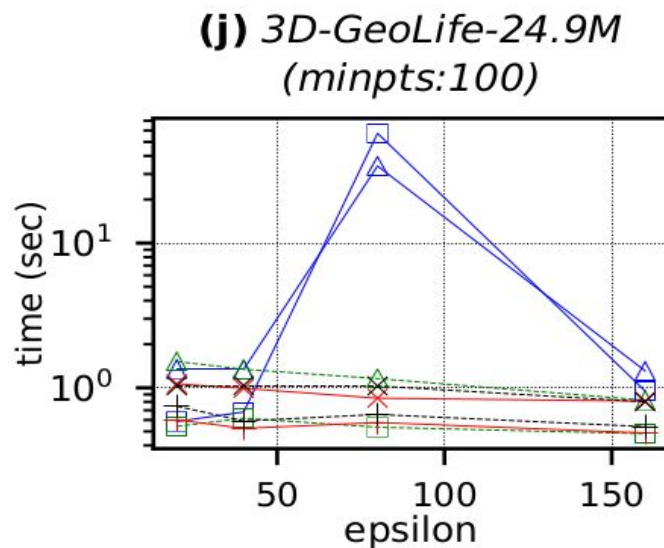
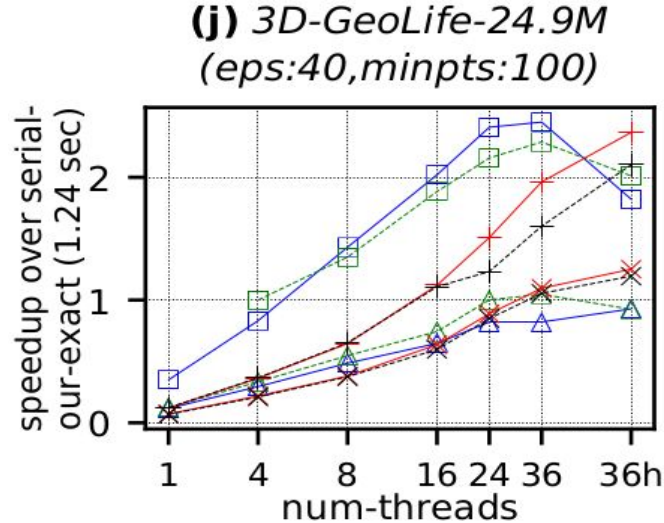
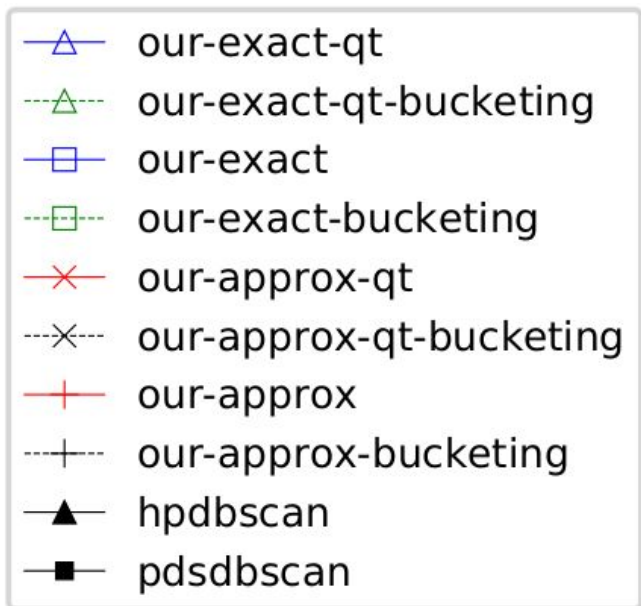
Experiments




Experiments



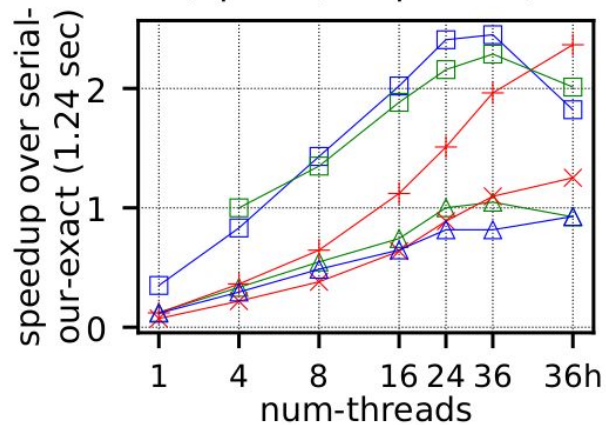
Experiments




Experiments



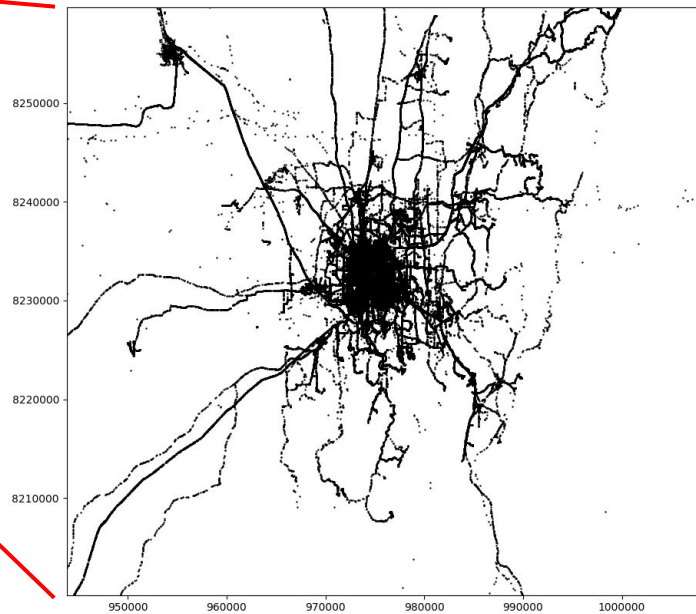
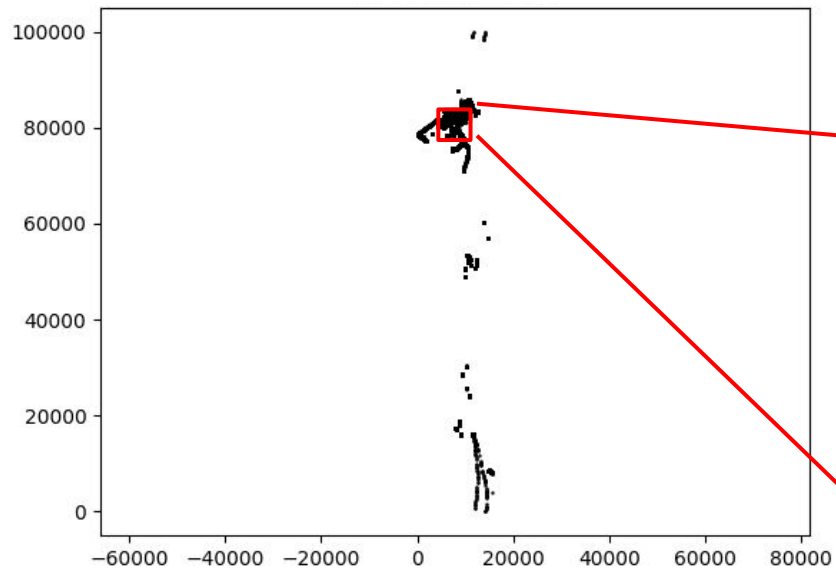
(j) 3D-GeoLife-24.9M
(eps:40,minpts:100)



(j) 3D-GeoLife-24.9M
(minpts:100)



GeoLife



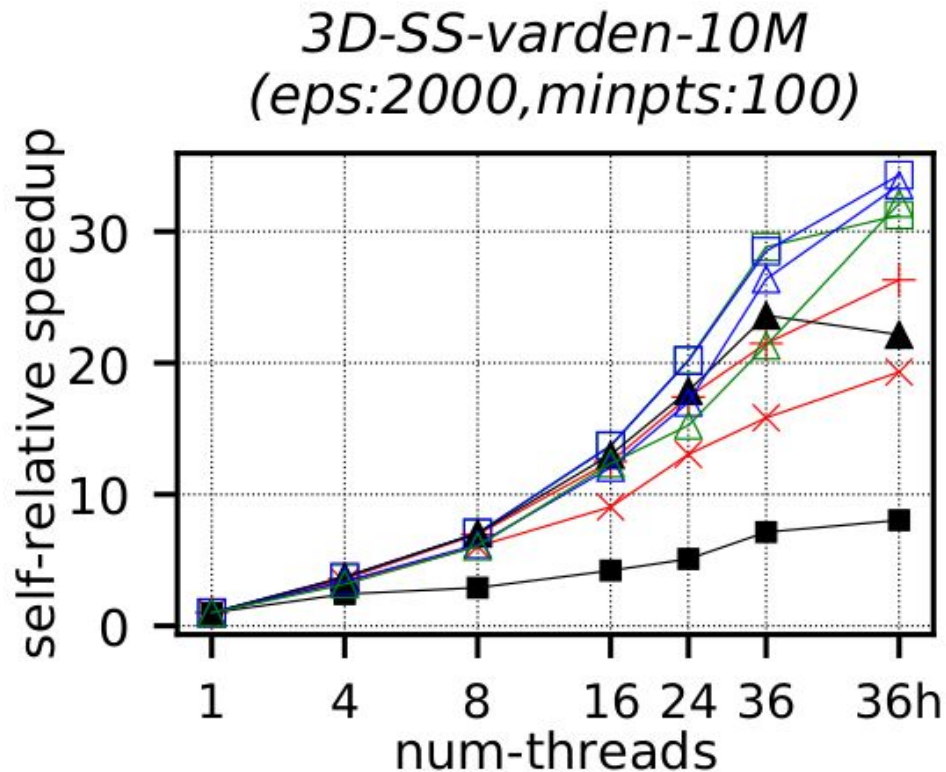

Experiments

	<i>#data</i>	<i>dimension</i>
<i>GeoLife</i>	<i>24.9M</i>	<i>3</i>
<i>Cosmo50</i>	<i>321M</i>	<i>3</i>
<i>OpenStreetMap</i>	<i>2770M</i>	<i>2</i>
<i>TeraClickLog</i>	<i>4373M</i>	<i>13</i>

	GeoLife				Cosmo50			
ϵ	20	40	80	160	0.01	0.02	0.04	0.08
<i>our-exact</i>	0.541	0.617	0.535	0.482	41.8	5.51	4.69	3.03
<i>rpdbscan (our machine)</i>	29.13	27.92	32.04	27.81	3750	562.0	576.9	672.6
<i>rpdbscan ([74])</i>	36	33	28	27	960	504	438	432

	OpenStreetMap				TeraClickLog			
ϵ	0.01	0.02	0.04	0.08	1500	3000	6000	12000
<i>our-exact</i>	41.4	43.2	40	44.5	26.8	26.9	27.0	27.6
<i>rpdbscan (our machine)</i>	–	–	–	–	–	–	–	–
<i>rpdbscan ([74])</i>	3000	1720	1200	840	15480	7200	3540	1680

Experiments - Self-relative speedups



Conclusion

- Practical parallel algorithms for 2D exact DBSCAN, and higher dimensional exact and approximate DBSCAN with work bounds matching the best sequential algorithms, and polylogarithmic depth.
- Highly-optimized implementations.
- A comprehensive experimental evaluation
 - 2-38x self-relative speedup (36 cores)
 - 5-33x speedup over best sequential (36 cores)
 - 14-6102x faster than existing parallel implementations
 - Processes largest dataset known for DBSCAN (> 4 billion points) on just one machine

Future work

- Our code will be available online soon

Questions