


# Massively Parallel Algorithms and Hardness of Single-Linkage Clustering

Joint work with Adithya Vadapalli (Indiana University)

**Grigory Yaroslavltssev**

<http://grigory.us>



# Algorithms for Big Data







- **User's perspective:** paradigm shift brought by cloud services
  - Outsourcing computation and data storage is great for both businesses and **researchers**
  - **Cloud service providers:** Amazon EC2, Google Compute Engine, Microsoft Azure, ...
  - **Open source stacks/frameworks:** MapReduce/Hadoop, Apache Spark, etc.





# Business perspective

- Pricings:
  - <https://cloud.google.com/pricing/>
  - <https://aws.amazon.com/pricing/>
- ~Linear with **space** and **time** usage
  - 100 machines: 5K \$/year
  - 10000 machines: 0.5M \$/year
- You pay **a lot more** for using provided algorithms
  - <https://aws.amazon.com/machine-learning/pricing/>

Compute Engine	
100 x	 
73,000 total hours per month	
VM class: regular	
Instance type: f1-micro	
Region: United States	
Sustained Use Discount: 30% ?	
Effective Hourly Rate: \$0.0056	
Estimated Component Cost: \$4,905.60 per 1 year	
1000 x	 
730,000 total hours per month	
VM class: regular	
Instance type: f1-micro	
Region: United States	
Sustained Use Discount: 30% ?	
Effective Hourly Rate: \$0.0056	
Estimated Component Cost: \$49,056.00 per 1 year	
10000 x	 
7,300,000 total hours per month	
VM class: regular	
Instance type: f1-micro	
Region: United States	
Sustained Use Discount: 30% ?	
Effective Hourly Rate: \$0.0056	
Estimated Component Cost: \$490,560.00 per 1 year	

# “Big Data Theory” = Turing meets Shannon




$$\begin{array}{c} \text{CPU time /} \\ \text{Computational} \\ \text{Complexity} \\ \\ = \quad + \\ \\ \text{Network Time /} \\ \text{Information and} \\ \text{Communication} \\ \text{Complexity} \end{array}$$



# Computational Model

- **Input:** size  $n$
- $M$  machines, space  $S$  on each ( $S = n^{1-\epsilon}$ ,  $0 < \epsilon < 1$ )
  - Constant overhead in total space:  $M \cdot S = O(n)$
- **Output:** solution to a problem (often size  $O(n)$ )
  - Doesn't fit on a single machine ( $S \ll n$ )




# Computational Model

- Computation/Communication in  $R$  rounds:
  - Every machine performs a **near-linear time** computation  $\Rightarrow$  Total running time  $O(S^{1+o(1)}R)$
  - Every machine **sends/receives at most  $S$  bits** of information  $\Rightarrow$  Total communication  $O(nR)$ .

**Goal:** Minimize  $R$ .

**Ideally:**  $R = \text{constant}$ .



# Algorithms for Graphs


- **Dense graphs vs. sparse graphs**

- **Dense:**  $S \gg |V|$

- Linear sketching: one round
    - “Filtering” [Karloff, Suri Vassilvitskii, SODA’10; Ene, Im, Moseley, KDD’11; Lattanzi, Moseley, Suri, Vassilvitskii, SPAA’11; Suri, Vassilvitskii, WWW’11] ... [Bateni, Behnezhad, Derakshan, Hajiaghayi, Kiveris, Lattanzi, Mirrokni, NIPS’17]

- **Sparse:**  $S \ll |V|$  (or  $S \ll$  solution size)

Sparse graph problems appear hard (**Big open question:** connectivity in  $o(\log n)$  rounds?)



# Two Conjectures

- **Conj 1: (Connectivity)** Given graph with  $O(|V|)$  edges finding connected components requires  $\Omega(\log |V|)$  rounds
- **Conj 2: (Two cycles)** Distinguishing one cycles from two cycles requires  $\Omega(\log |V|)$  rounds
- [Roughgarden, Vassilvitskii, Wang SPAA'16]
  - Give an  $\Omega(\log_S n)$  lower bound
  - Some evidence of relationship to  $NC^1 \subseteq P$




# Geometric Graph Problems [ANOE'14]

Combinatorial problems on graphs in  $\mathbb{R}^d$

## Polynomial time (“easy”)

- Minimum Spanning Tree
- Earth-Mover Distance =  
Min Weight Bi-chromatic Matching




## ~~NP-hard (“hard”)~~

- ~~• Steiner Tree~~
- ~~• Traveling Salesman~~
- ~~• Clustering (k-medians, facility location, etc.)~~

~~Arora-Mitchell-style  
“Divide and Conquer”,  
easy to implement in  
Massively Parallel  
Computational Models,  
but bad running time~~

# MST: Single Linkage Clustering


- [Zahn'71] **Clustering** via MST (Single-linkage):  
 $k$  clusters: remove  $k - 1$  longest edges from MST
- Maximizes **minimum** intercluster distance



Objective:  $\min(a, b, c)$

# MST vs. Single Linkage Clustering

- [ANOY'14]:  $(1 + \epsilon)$ -approx MST (in expectation)
- Single-linkage clustering can be arbitrarily bad:



# Our Results

	Approximation in $O(\log n)$ rounds	Hardness of approx. in $o(\log n)$ rounds
$\ell_0$	Exact for $d = O(1)$	2 for $d = 2$ under Connectivity 3 for $d = \Omega(n)$ under Two Cycles
$\ell_1$	$(1 + \epsilon)$ for $d = O(1)$	2 for $d = \Omega(n)$ under Connectivity 3 for $d = \Omega(n)$ under Two Cycles
$\ell_2$	$(1 + \epsilon)$ for $d = O(1)$	$1.41 - \epsilon$ for $d = \Omega\left(\frac{\log n}{\epsilon^2}\right)$ under Connectivity $1.84 - \epsilon$ for $d = \Omega\left(\frac{\log n}{\epsilon^2}\right)$ under Two Cycles
$\ell_\infty$	$(1 + \epsilon)$ for $d = O(1)$	2 for $d = \Omega(n)$ under Connectivity [ANOY'14]

$\ell_0$  and  $\ell_1$  hardness results hold for  $O(1)$ -sparse vectors, i.e. for inputs of size  $O(n)$

# Hardness for $\ell_2$

## Reduction from “One vs. Two Cycles”

- Vector  $v_i$  for each vertex, set  $v_i = e_i$
- For each edge  $(i, j)$  update (for  $\xi = 1/\sqrt{2}$ ):
  - $v_i = v_i + \xi e_j$
  - $v_j = v_j + \xi e_i$
- Apply Johnson-Lindenstrauss transform to reduce the dimension down to  $d = O\left(\frac{\log n}{\epsilon^2}\right)$

**Important that reduction can be done in  $O(1)$  rounds**

# Hardness for $\ell_2$

## Reduction from “One vs. Two Cycles”

- Vector  $\mathbf{v}_i$  for each vertex, set  $\mathbf{v}_i = \mathbf{e}_i$
- For each edge  $(i, j)$  update (for  $\xi = 1/\sqrt{2}$ ):
  - $\mathbf{v}_i = \mathbf{v}_i + \xi \mathbf{e}_j$
  - $\mathbf{v}_j = \mathbf{v}_j + \xi \mathbf{e}_i$
- $\|\mathbf{v}_i - \mathbf{v}_j\|_2 = \sqrt{2} (\sqrt{2} - \sqrt{2})$  if there is an edge  $(i, j)$
- $\|\mathbf{v}_i - \mathbf{v}_j\|_2 = 2$  if there is no edge  $(i, j)$

Ratio between these cases gives hardness of  $\sqrt{2 + \sqrt{2}}$

# General algorithm for $\ell_1, \ell_2, \ell_\infty$

- Input: vectors  $v_1, \dots, v_n \in \mathbb{R}^d$
- $E = \emptyset$
- Repeat  $O(\log n)$  times sequentially:
  - $E' =$  set of edges of a  $(1 + \epsilon)$ -approximate MST
  - $E = E \cup E'$
- Run Boruvka's algorithm on  $E$
- Drop  $k - 1$  longest edges to get the clustering

# Large geometric graphs

- Graph algorithms: **Dense graphs** vs. sparse graphs
  - **Dense:**  $S \gg |V|$ .
  - **Sparse:**  $S \ll |V|$ .
- Our setting:
  - Dense graphs, sparsely represented:  $O(n)$  space
  - Output doesn't fit on one machine ( $S \ll n$ )
- **Today:**  $(1 + \epsilon)$ -approximate MST [ANOV'14]
  - $d = 2$  (easy to generalize)
  - $R = \log_S n = O(1)$  rounds ( $S = n^{\Omega(1)}$ )




# $O(\log n)$ -MST in $R = O(\log n)$ rounds

- Assume points have integer coordinates  $[0, \dots, \Delta]$ , where  $\Delta = O(n^2)$ .

Impose an  $O(\log n)$ -depth quadtree

Bottom-up: For each cell in the quadtree

- compute optimum MSTs in subcells
- Use only **one representative** from each cell on the next level



**Wrong representative:**

$O(1)$ -approximation per level



# $\epsilon L$ -nets

- $\epsilon L$ -net for a cell  $C$  with side length  $L$ :  
Collection  $S$  of vertices in  $C$ , every vertex is at distance  $\leq \epsilon L$  from some vertex in  $S$ . (Fact: Can efficiently compute  $\epsilon$ -net of size  $O\left(\frac{1}{\epsilon^2}\right)$ )

Bottom-up: For each cell in the quadtree

- Compute optimum MSTs in subcells
- Use  $\epsilon L$ -net from each cell on the next level

- **Idea:** Pay only  $O(\epsilon L)$  for an **edge** cut by cell with side  $L$
- Randomly shift the quadtree:  
 $\Pr[\text{cut edge of length } \ell \text{ by } L] \sim \ell/L$  – charge errors  $O(1)$ -approximation per level




# Randomly shifted quadtree

- Top cell shifted by a random vector in  $[0, L]^2$

Impose a **randomly shifted** quadtree (top cell length  $2\Delta$ )

Bottom-up: For each cell in the quadtree

- Compute optimum MSTs in subcells
- Use  $\epsilon L$ -net from each cell on the next level



Pay **5** instead of **4**  
**Bad Cut**  
 $\Pr[\text{Bad Cut}] = \Omega(1)$


# $(1 + \epsilon)$ -MST in $\mathbf{R} = O(\log n)$ rounds

- **Idea:** Only use short edges inside the cells

Impose a **randomly shifted** quadtree (top cell length  $\frac{2\Delta}{\epsilon}$ )

Bottom-up: For each node (cell) in the quadtree

- compute optimum Minimum Spanning **Forests** in subcells, **using edges of length  $\leq \epsilon L$**
- Use only  $\epsilon^2 L$ -net from each cell on the next level




$$L = \Omega\left(\frac{1}{\epsilon}\right)$$

$$\Pr[\mathbf{Bad\ Cut}] = O(\epsilon)$$

# $(1 + \epsilon)$ -MST in $\mathbf{R} = O(1)$ rounds

- $O(\log n)$  rounds  $\Rightarrow O(\log_s n) = O(1)$  rounds
  - Flatten the tree:  $(\sqrt{M} \times \sqrt{M})$ -grids instead of  $(2 \times 2)$  grids at each level.



Impose a **randomly shifted**  $(\sqrt{M} \times \sqrt{M})$ -tree

Bottom-up: For each node (cell) in the tree

- compute optimum MSTs in subcells via edges of length  $\leq \epsilon L$
- Use only  $\epsilon^2 L$ -net from each cell on the next level

# $(1 + \epsilon)$ -MST in $R = O(1)$ rounds

**Theorem:** Let  $l = \#$  levels in a random tree  $P$

$$\mathbb{E}_P[\mathbf{ALG}] \leq (1 + O(\epsilon l d)) \mathbf{OPT}$$

**Proof (sketch):**

- $\Delta_P(u, v)$  = cell length, which first partitions  $(u, v)$
- **New weights:**  $w_P(u, v) = ||u - v||_2 + \epsilon \Delta_P(u, v)$

$$||u - v||_2 \leq \mathbb{E}_P[w_P(u, v)] \leq (1 + O(\epsilon l d)) ||u - v||_2$$

- Our algorithm implements Kruskal for weights  $w_P$

# Thank you!

- Experiments in Apache Spark on largest vector datasets from UCI ML repository
  - $\approx 11\text{M}$  vectors  $\Rightarrow$  960 TB for adjacency matrix
  - SIFT & HIGGS datasets preprocessed with PCA
- More on my blog <http://grigory.us/blog/>
- CAML: <http://caml.indiana.edu>