

Nimble Algorithms for Cloud Computing

Ravi Kannan, Santosh Vempala and David Woodruff

Cloud computing

Data is distributed arbitrarily on many servers

Parallel algorithms: time

Streaming algorithms: sublinear space

Cloud Complexity: time, space *and communication*

[Cormode-Muthukrishnan-Ye 2008]

Nimble algorithm: polynomial time/space (as usual) and sublinear (ideally polylog) communication between servers.



Cloud vs Streaming

- ▶ Streaming algorithms make small “sketches” of data
- ▶ Nimble algorithms must communicate small “sketches”
- ▶ Are they equivalent?

Simple observation:

- ▶ Communication in cloud = $O(\text{memory in streaming})$

[Daume-Philips-Saha-Venkatasubramanian 12]

- ▶ Is cloud computing more powerful?
-



Basic Problems on large data sets

- ▶ Frequency moments
- ▶ Counting copies of subgraphs (homomorphisms)
- ▶ Low-rank approximation
- ▶ Clustering

- ▶ ...
- ▶ Matchings
- ▶ Flows
- ▶ Linear programs



Streaming Lower Bounds

Frequency moments: Given a vector of frequencies $f = (f_1, f_2, \dots, f_n)$ presented as a set of increments, estimate $\|f\|_k = \sum_{i=1}^n f_i^k$ to relative error ϵ .

[Alon-Matias-Szegedy99, Indyk-Woodruff05]:

$\Theta(n^{1-2/k})$ space (k = 1, 2 by random projection)

Counting homomorphisms: Estimate #triangles, #C4, #K_r, ... in a large graph G.

$\Omega(n^2)$ space lower bounds in streaming.



Streaming Lower Bounds

Low-rank approximation: Given $n \times d$ matrix A , find A of rank k
s.t.

$$\|A - A\|_F \leq (1 + \epsilon) \|A - A_k\|_F$$

[Clarkson-Woodruff09]

Any streaming algorithm needs $\Omega((n+d)k \log nd)$ space.



Frequency moments in the cloud

- ▶ Lower bound via multi-player set disjointness.
- ▶ t players have sets S_1, S_2, \dots, S_t , subsets of $[n]$
- ▶ Problem: determine if sets are disjoint or have one element in common.
- ▶ Thm: Communication needed = $\Omega(n/t \log t)$ bits.



Frequency moments in the cloud

Thm. Communication needed to determine set disjointness of t sets is $\Omega(n/t \log t)$ bits.

Consider s sets being either

(i) completely disjoint or (ii) with one common element
(each set is on one server)

Then k 'th frequency moment is either n or $n-1+s^k$

Suppose we have a factor 2 approximation for the k 'th moment. With $s^k = n+1$, then we can distinguish these cases. Therefore, communication needed is $\Omega(s^k - 1)$.



Frequency moments in the cloud

Thm. [Kannan-V.-Woodruff13]


Estimating k 'th frequency moment on s servers takes $O(s^k / \epsilon^2)$ words of communication, with $O(b + \log n)$ bits per word.

- ▶ Lower bound is s^{k-1}
- ▶ Previous bound: $s^{k-1} (\log n / \epsilon)^{O(k)}$ [Woodruff-Zhang12]
- ▶ streaming space complexity is $n^{1-2/k}$

Main idea of algorithm: sample elements within a server according to higher moments.



Warm-up: 2 servers, third moment



Goal: estimate $\sum_{i=1}^n (u_i + v_i)^3$

1. Estimate $\sum_{i=1}^n u_i^3$
2. Sample j w.p. $p_j = u_j^3 / \sum_{i=1}^n u_i^3$;
announce
3. Second server computes $X = u_j^2 v_j / p_j$
4. Average over many samples.

$$E(X) = \sum_{i=1}^n u_i^2 v_i$$



Warm-up: 2 servers, third moment

Goal: estimate $\sum_{i=1}^n (u_i + v_i)^3$

$$p_j = u_j^3 / \sum_{i=1}^n u_i^3 \quad X = u_j^2 v_j / p_j \quad E(X) = \sum_{i=1}^n u_i^2 v_i$$

$$\text{Var}(X) \leq \sum_{i: v_i > 0} (u_i^2 v_i)^2 / p_i$$

$$\leq \sum_{i=1}^n u_i^3 \sum_{i=1}^n u_i v_i^2$$

$$\leq (\sum_{i=1}^n u_i^3 + v_i^3)^2$$

So, $O(1/\epsilon^2)$ samples suffice.



Many servers, k'th moment

GOAL: $\sum_{i=1}^n \left(\sum_{j=1}^p f_{ij} \right)^k$

$$= \sum_i \sum_{\substack{v_1, \dots, v_p \\ \sum v_j = k}} \binom{k}{v_1, v_2, \dots, v_p} \prod_j f_{ij}^{v_j}$$

$$= \sum_{v_1, \dots, v_p} \binom{k}{v_1, \dots, v_p} \underbrace{\sum_i \prod_j f_{ij}^{v_j}}$$



Many servers, k'th moment

GOAL:
$$\sum_i f_{i_1}^{y_1} \cdots f_{i_m}^{y_m} \quad (\sum y_j = k)$$

Each server j:

- ▶ Sample i w. prob $p_i = f_{ij}^k / \sum_t f_{tj}^k$ according to k'th moment.
 - ▶ Every j' sends $f_{ij'}^k$ if $j' < j$ and $f_{ij'}^k < f_{ij}^k$
or $j' > j$ and $f_{ij'}^k \leq f_{ij}^k$
 - ▶ Server j computes $X_i = \prod_{j=1}^m f_{ij}^{y_j} / p_i$
-



Many servers, k'th moment

Each server j:

- ▶ Sample i w. prob $p_i = f_{ij}^k / \sum_t f_{tj}^k$ according to k'th moment.
- ▶ Every j' sends $f_{ij'}^k$ if $j' < j$ and $f_{ij'}^k < f_{ij}^k$
or $j' > j$ and $f_{ij'}^k \leq f_{ij}^k$
- ▶ Server j computes $X_i = \prod_{j=1}^s f_{ij}^k r_j / p_i$

Lemma. $E(X) = \sum R_j \prod_{j=1}^s f_{ij}^k r_j$ and $Var(X) \leq (\sum_i f_{ij}^k)^2$

Theorem follows as there are $< s^k$ terms in total.



Counting homomorphisms

- ▶ How many copies of graph H in large graph G ?
- ▶ E.g., $H =$ triangle, 4-cycle, complete bipartite etc.
- ▶ Linear lower bounds for counting 4-cycles, triangles.
- ▶ We assume an (arbitrary) partition of the vertices among servers.



Counting homomorphisms

- ▶ To count number of paths of length 2, in a graph with degrees d_1, d_2, \dots, d_n , we need:

$$t(K_{1,2}, G) = \sum_{i=1}^n d_i^2$$

This is a polynomial in frequency moments!

- ▶ #stars is $t(K_{1,r}, G) = \sum_{i=1}^n d_i^r$
- ▶ #C4's: let d_{ij} is the number of common neighbors of i and j . Then,

$$t(C_4, G) = \sum_{i,j} d_{ij}^2$$

- ▶ # $K_{a,b}$: let d_S be the number of common neighbors of a set of vertices S . Then,

$$t(K_{a,b}, G) = \sum_{S \subset V, |S|=a} d_S^b$$



Low-rank approximation

Given $n \times d$ matrix A partitioned arbitrarily as

$A = A \downarrow 1 + A \downarrow 2 + \dots + A \downarrow s$ among s servers, find A of rank k s.t.

$$\|A - A\|_F \leq (1 + \epsilon) OPT.$$


To avoid linear communication, on each server t , we leave a matrix $A \downarrow t$, s.t. $A = A \downarrow 1 + A \downarrow 2 + \dots + A \downarrow s$ and is of rank k .

How to compute these matrices?



Low-rank approximation in the cloud

Thm. [KVW13]. Low-rank approximation of $n \times d$ matrix A partitioned arbitrarily among s servers takes $O(\uparrow^* (skd))$ communication.



Warm-up: row partition

- ▶ Full matrix A is $n \times d$ with $n \gg d$.
- ▶ Each server j has a subset of rows $A_{\downarrow j}$
- ▶ Computes $A_{\downarrow j} A_{\downarrow j}^T$ and sends to server 1.
- ▶ Server 1 computes $B = \sum_{j=1}^s A_{\downarrow j} A_{\downarrow j}^T$ and announces V , the top k eigenvectors of B .
- ▶ Now each server j can compute $A_{\downarrow j} V V^T$.
- ▶ Total communication = $O(sd^2)$.



Low-rank approximation: arbitrary partition

- ▶ To extend this to arbitrary partitions, we use limited-independence random projection.
- ▶ Subspace embedding: matrix P of size $O(d/\epsilon^2) \times n$ s.t. for any $x \in \mathbb{R}^d$, $\|PAx\| = (1 \pm \epsilon)\|Ax\|$.
- ▶ Agree on projection P via a random seed
- ▶ Each server computes $PA_{\downarrow t}$, sends to server 1.
- ▶ Server 1 computes $PA = \sum_t \uparrow PA_{\downarrow t}$ and its top k right singular vectors V .
- ▶ Project rows of A to V .
- ▶ Total communication = $O(sd^2/\epsilon^2)$.



Low-rank approximation: arbitrary partition

- ▶ Agree on projection P via a random seed
- ▶ Each server computes $PA \downarrow t$, sends to server 1.
- ▶ Server 1 computes $PA = \sum_{t=1}^T PA \downarrow t$ and its top k right singular vectors V .
- ▶ Project rows of A to V .

Thm. $\|A - AVV^T\| \leq (1 + O(\epsilon))OPT$.

Pf. Extend V to a basis $v \downarrow 1, v \downarrow 2, \dots, v \downarrow d$. Then,

$$\|A - AVV^T\|_{F^2} = \sum_{i=k+1}^d \|Av \downarrow i\|^2 \leq (1 + \epsilon)^2 \sum_{i=k+1}^d \|PAv \downarrow i\|^2 .$$

And, with $u \downarrow 1, u \downarrow 2, \dots, u \downarrow d$ singular vectors of A ,

$$\begin{aligned} \sum_{i=k+1}^d \|PAv \downarrow i\|^2 &\leq \sum_{i=k+1}^d \|PAu \downarrow i\|^2 \leq (1 + \epsilon)^2 \sum_{i=k+1}^d \|Au \downarrow i\|^2 \\ &= (1 + O(\epsilon))OPT^2 . \end{aligned}$$



Low-rank approximation in the cloud

To improve to $O(kd)$, we use a subspace embedding up front, and observe that $O(k)$ -wise independence suffices for the random projection matrix.

- ▶ Agree on $O(k/\epsilon) \times n$ matrix S and $O(k/\epsilon^2) \times n$ matrix P .
- ▶ Each server computes $SA \downarrow t$ and sends to server 1.
- ▶ S1 computes $SA = \sum_t \uparrow \text{matrix} \downarrow SA \downarrow t$ and an orthonormal basis $U \uparrow T$ for its row space.
- ▶ Apply previous algorithm to AU .



K-means clustering

- ▶ Find a set of k centers c_1, c_2, \dots, c_k that minimize $\sum_{i \in S} \min_{j=1 \dots k} \|x_i - c_j\|^2$
- ▶ A near-optimal (i.e. $1 + \epsilon$) solution could be very different!
- ▶ So, cannot project up front to reduce dimension and approximately preserve distances.



K-means clustering

- ▶ **Kannan-Kumar condition:**
- ▶ Every pair of cluster centers are $f(k)$ standard deviations apart.
- ▶ “variance”: maximum over 1-d projections, of the average squared distance of a point to its center.
(e.g. for Gaussian mixtures, max directional variance)
- ▶ Thm. [Kannan-Kumar 10]. Under this condition, projection to the top k principal components followed by the k -means iteration starting at an approximately optimal set of centers finds a nearly correct clustering.
- ▶ Finds centers close to the optimal ones, so that the induced clustering is same for most point.



K-means clustering in the cloud

- ▶ Points (rows) are partitioned among servers
- ▶ Low-rank approximation to project to SVD space.
- ▶ How to find a good starting set of centers?
- ▶ Need a constant-factor approximation.
- ▶ Thm [Chen]. There exists a small subset (“core”) s.t. the k-means value of this set (weighted) is within a constant factor of the k-means value of the full set of points (for any set of centers!).
- ▶ Chen’s algorithm can be made nimble.

Thm. K-means clustering in the cloud achieves the Kannan-Kumar guarantee with $O(d^2 + k^4)$ communication on $s = O(1)$ servers.



Cloud computing: What problems have nimble algorithms?

- ▶ Approximate flow/matching?
- ▶ Linear programs
- ▶ Which graph properties/parameters can be checked/estimated in the cloud?
(e.g., planarity? expansion? small diameter?)
- ▶ Other Optimization/Clustering/
- ▶ Learning problems
[Balcan-Blum-Fine-Mansour | 2,
Daume-Philips-Saha-Venkatasubramaian | 2]
- ▶ Random partition of data?
- ▶ Connection to property testing?

