# Lectures on Streaming Algorithms Qin Zhang 

May 13, 2011

## 1 Models and Motivations

### 1.1 Models

- Standard stream model: $m$ elements from universe of size $n$, come one by one.

Goal: compute a function of stream.
Constraints: (1) Limited space (working memory), sublinear in $n$ and $m$.
(2) Access data sequentially. (3) Process each element quickly.

Two variants: (1) Cash register (2) Turnstile.

- Graph stream: elements of the stream are edges. Let $n$ be the number of vertices. We have $m=O\left(n^{2}\right)$. To be able to handle interesting functions, we typically allow space $O\left(n \log ^{O(1)} n\right)$.
- Sliding windows: (1) Sequence-based sliding window: only interested in the last $w$ elements. (2) Time-based sliding window: only interested in elements that arrive in the last $w$ time steps.


### 1.2 Motivations

- Practical Appeal: Network monitoring (\#distinct(sourse, destination) pairs, top- $k$ heaviest flows), query planning, I/O efficiency for massive data, sensor networks aggregation ....
- Theoretical Appeal: Easy to state problems but hard to solve. Links to communication complexity, compressed sensing, embeddings, pseudo-random generators, approximation ....


## 2 Samplings

### 2.1 Motivation

Sampling is a general technique for tackling massive amounts of data. For example, to compute the median packet size of some IP packets, we could just sample some and use the median of the sample as an estimate for the true median. Statistical arguments relate the size of the sample to the accuracy of the estimate.

### 2.2 Standard Sampling

- Q: Sampling from a stream of unknown length
- A: Reservoir Sampling

1. Initially $s=x_{1}$
2. On seeing the $t$-th element, set $s$ to be $x_{t}$ with probability $1 / t$.

- Q: Sampling from a sliding window (last $w$ items).
- A: Algorithm:

1. For each $x_{i}$ we pick a random value $v_{i} \in(0,1)$.
2. In a window $\left\langle x_{j-w+1} ; \ldots ; x_{j}\right\rangle$ return value $x_{i}$ with smallest $v_{i}$.

How? Maintain a set of all elements in the sliding window whose value is minimal among subsequent values.

Cost: $O(\log n \log w)$ space expected.

- Q: Can we make space optimal $(O(\log n))$ and worst case?
- A: Yes. See Optimal sampling from sliding windows by Braverman et. al. PODS 2009.


### 2.3 AMS Sampling

- Q: Want to compute some function $\sum_{i \in[n]} g\left(f_{i}\right)$ where $f_{i}=\left|\left\{j: x_{j}=i\right\}\right|$ is the frequency of $i$ and $g$ is an arbitrary function with $g(0)=0$.
- A: Algorithm:

1. Pick $J$ uniformly at random from $[m]$.
2. Compute $R=\left|\left\{j \geq J: a_{j}=a_{J}\right\}\right|$.
3. Let $X=m(g(R)-g(R-1))$

Cost: $O(\log m+\log n)$ space (1-run) guarantees $E[X]=\sum_{i \in[n]} g\left(f_{i}\right)$.

### 2.3.1 Analysis

$$
\begin{aligned}
\mathbf{E}[X] & =\sum_{i} \mathbf{E}\left[X \mid a_{J}=i\right] \operatorname{Pr}\left[a_{J}=i\right] \\
& =\sum_{i} \mathbf{E}\left[g(R)-g(R-1) \mid a_{J}=i\right] f_{i} \\
& =\sum_{i}\left(\frac{g(1)-g(0)}{f_{i}}+\ldots+\frac{g\left(f_{i}\right)-g\left(f_{i}-1\right)}{f_{i}}\right) f_{i} \\
& =\sum_{i} g\left(f_{i}\right)
\end{aligned}
$$

## 3 Heavy Hitters

- Q: Let $f_{i}=\left|\left\{j: a_{j}=i\right\}\right|$. How to maintain for each $i \in[n]$ an estimate $\tilde{f}_{i}$ s.t. $f_{i} \leq \tilde{f}_{i} \leq f_{i}+\epsilon m$ ?
- A: Algorithm [CR-Precis]: Let $p_{1}, p_{2}, \ldots, p_{t}$ be the first $t$ prime numbers and define hash functions $h_{j}(x)=x \bmod p_{j}$. Next, maintain $p_{t} t=O\left(t^{2} \log t\right)$ counters $c_{i, j}$ which are initially 0 . When a new item $e$ comes, increment each of $c_{j, h_{j}(e)}$ for $j \in[t]$.
At the end, we set $\tilde{f}_{i}=\min _{j} c_{j, h_{j}(i)}$.
Cost: set $t=O\left(\epsilon^{-1} \log n\right)$ and the space cost is $O\left(\epsilon^{-2} \log ^{2} n(\log 1 / \epsilon+\right.$ $\log \log n)$ ).
- Q: Can we do better?
- A: Sure. Algorithm[Space-saving, by Metwally et. al. An integrated efficient solution for computing frequent and top-k elements in data streams. TODS 2006]: We construct an array with $1 / \epsilon$ cells. Each cell is of the form $(e, f(e))$, where $f(e)$ denotes the estimated frequency of the item $e$.
When a new item $e$ comes, we have two cases.

1. If $e$ is not in the array, we create a new tuple $(e, \operatorname{MIN}+1)$ where $\operatorname{MIN}=\min \left\{f\left(e_{i}\right): e_{i}\right.$ is in the array $\}$. We always keep the array sorted according to $f\left(e_{i}\right)$, and then MIN is just the estimated frequency of the last item. If the length array is larger than $1 / \epsilon$, we delete the last tuple.
2. If $e$ is already in the array. We just increment $f(e)$ by 1 and reinsert the $(e, f(e))$ into the array.

Cost: $O(1 / \epsilon)$.

- A: So what is the merit of the first (cumbersome) algorithm?
- Q: It is useful since it is an linear mapping; it can handle deletions.

Linear mapping is one technique to perform sketching. The basic idea of sketching is to apply a linear projection "on the fly" that takes high-dimensional data to a smaller dimensional space. Post-process lower dimensional image to estimate the quantities of interest. Since it is a linear mapping, it can support "deletions".

### 3.1 Analysis of CR-Precis

For any $i^{\prime} \neq i$, there are at most $\log n$ functions $h_{j}$ under which $i, i^{\prime}$ collide, by the Chinese Remainder Theorem. Hence, we deduce that,

$$
\sum_{j} c_{j, h_{j}(i)} \leq t f_{i}+\left(\sum_{i^{\prime} \neq i} f_{i^{\prime}}\right) \log n
$$

Thus $\tilde{f}_{i} \leq f_{i}+(m \log n) / t$.

## 4 Distinct Elements

- Q: Compute $d=\left|\left\{j: f_{j}>0\right\}\right|$. The number of distinct elements.
- A: Algorithms [FM, Flajoet and Martin, FOCS 1983]: let $\operatorname{zeros}(p)=$ $\max \left\{i: 2^{i}\right.$ divides $\left.p\right\}$.

1. Choose a random hash function $h:[n] \rightarrow[n]$ from a 2-universal family. Set $z=0$.
2. For each new coming item $e$, if $\operatorname{zeros}(h(e))>z$, then $\operatorname{set} z=\operatorname{zeros}(h(e))$;
3. Output $2^{z+0.5}$.

Cost: we should repeat it $\log (1 / \delta)$ times (i.e. maintain $\log (1 / \delta)$ hash functions in parallel) and then take the median to get a successful probability of $1-\delta$. Thus need $O(\log (1 / \delta) \log n)$ space.

- Q: Can we do better in terms of approximation ratio?
- A: Yes. Algorithm [BJKST]

1. Choose a random hash function $h:[n] \rightarrow[n]$ from a 2-universal family. Set $z=0, B=\emptyset$. Choose a secondary 2-universal hash function $g:[n] \rightarrow\left[O\left(\log n / \epsilon^{2}\right)\right]$.
2. For each new coming item $e$, if $\operatorname{zeros}(h(e)) \geq z$, then
(a) set $B=B \cup\{g(e)\}$;
(b) if $|B|>c / \epsilon^{2}$ then set $z=z+1$ and rehash $B$.
3. Output $|B| 2^{z}$.

Cost: $O\left(\log n+1 / \epsilon^{2} \cdot(\log 1 / \epsilon)+\log \log n\right)$.

- Q: Can we do better?
- A: Yes. See the optimal algorithm in An Optimal Algorithm for the Distinct Elements Problem by Kane et. al. PODS 2010.


### 4.1 Analysis

[FM] Basic intuition is that we expect 1 out of the $d$ distinct elements to hit $\operatorname{zeros}(h(j))>\log d$, and we don't expect any elements to hit $\operatorname{zeros}(h(j)) \gg \log d$. For details, see [LN], page 9 .

The median trick. Suppose we know that for a random variable $X$, with probability no more than $1 / 3$ (could be replaced by any constant $<1 / 2$ ), $X \leq a$; and with probability no more than $1 / 3, X \geq b$. Now if we run $k$ copies of the algorithm in parallel, thus get $k X_{i} \mathrm{~s}$, and then take the median of $X_{i} \mathrm{~s}$, say, $\bar{X}$. We have that $\bar{X} \in[a, b]$ with probability at least $1-2^{-\Omega(k)}$.
[BJKST] The algorithm is in the same spirit of [FM], except that at the high end of the "FM hit-array", they introduced a "buffer" to get a finer estimation. For details, see [LN], pages 17-18.

## 5 AMS Estimator for $F_{k}$

- Q: How to compute $F_{k}=\sum_{j}\left|f_{j}\right|^{k}$ where $f_{j}$ is the frequency of item $j$ in the stream?
- A: Use AMS sampling introduced before (Homework).
- Q: Better bounds?
- A: Yes. See the optimal algorithm in Optimal Approximations of the Frequency Moments of Data Streams by Indyk and Woodruff. STOC 2005.
- Q: Can we do even faster for $F_{2}$ ?
- A: Yes. Algorithm: let $f_{i}=\left|\left\{j: a_{j}=i\right\}\right|$.

1. For $1 \leq i \leq s_{1}=\frac{16}{\lambda^{2}}$ and $1 \leq j \leq s_{2}=2 \log (1 / \epsilon)$, generate $X_{i, j}$ as follows:
(a) Choose a random $\{-1,+1\}$ vector $v=\left\{\epsilon_{1}, \ldots, \epsilon_{n}\right\}$ from a 4 -wise $\{-1,+1\}$ vectors family, compute $Z=\sum_{i} \epsilon_{i} f_{i}$.
(b) Let $X=Z^{2}$.
2. Let $Y_{j}=\sum_{i} X_{i, j} / s_{1}$
3. Output the median of $Y_{j} \mathrm{~s}$.

- Q: Why this works? Any general techniques?
- A: Yes. The $p$-stable distributions. A distribution is $p$-stable if $X_{1}, X_{2}, \ldots, X_{n} \sim$ $\mu$ i.i.d. and $\vec{a}=\left(a_{1}, \ldots, a_{n}\right) \in \mathbb{R}^{n}$, then $a_{1} X_{1}+a_{2} X_{2}+\ldots+a_{n} X_{n} \sim\|\vec{a}\|_{p} X$ where $X \sim \mu$.
For example, let $X_{1}, X_{2}, \ldots, X_{n} \sim N(0,1)$. Then we have $a_{1} X_{1}+a_{2} X_{2}+$ $\ldots+a_{n} X_{n} \sim \sqrt{a_{1}^{2}+a_{2}^{2}+\ldots+a_{n}^{2}} X$ where $X \sim N(0,1)$
Cauchy distribution is 1 -stable where the Cauchy pdf is $\frac{2}{\pi} \frac{1}{1+x^{2}}$.
For details, see P. Indyk Stable Distributions, Pseudo-random Generators, Embeddings and Data Stream Computation. STOC 2006.
- Q: Then why $\{+1,-1\}$ works?
- A: Because if $Y_{1}, \ldots, Y_{n} \in_{R}\{+1,-1\}$, then $\vec{Y}=\left\{Y_{1}, \ldots, Y_{n}\right\}$ is "close" to $\vec{Z}=\left\{Z_{1}, \ldots, Z_{n}\right\}$ where $Z_{i} \in N(0,1)$ for $n \rightarrow \infty$.


### 5.1 Analysis of AMS Estimator for $F_{2}$

Note that $\mathbf{E}\left[\epsilon_{i}\right]=0$ for all $i$.

$$
\begin{aligned}
\mathbf{E}[X] & =\mathbf{E}\left[\left(\sum_{i} \epsilon_{i} f_{i}\right)^{2}\right]=\sum_{i} f_{i}^{2} \mathbf{E}\left[\epsilon_{i}^{2}\right]+2 \sum_{i, j} f_{i} f_{j} \mathbf{E}\left[\epsilon_{i}\right] \mathbf{E}\left[\epsilon_{j}\right]=\sum_{i} m_{i}^{2}=F_{2} . \\
\operatorname{Var}[X] & =E\left[X^{2}\right]-(E[X])^{2}=4 \sum_{i, j} m_{i}^{2}, m_{j}^{2} \leq 2 F_{2}^{2} .
\end{aligned}
$$

Now plug in average + median.

$$
\operatorname{Pr}\left[\left|Y_{i}-F_{2}\right|>\lambda F_{2}\right] \leq \frac{\operatorname{Var}\left(Y_{i}\right)}{\lambda^{2} F_{2}^{2}} \leq 1 / 8
$$

## 6 Graph Algorithms

### 6.1 Connectivity

- Q: Decide whether a graph is connected in small space.
- A: Algorithm:

1. Maintain label $\ell(u)$ for each node $u$ where labels are initially distinct
2. On seeing edge $(u ; v)$ with $\ell(u) \neq \ell(v)$, set $\ell(w)$ for all $\ell(u)$ for all $w$ with $\ell(w)=\ell(v)$.
3. The graph is connected iff every node ends up with the same label.

### 6.2 Test Bipartite Graphs

- Q: Decide whether a graph is connected or not in $O(n \log n)$ space.
- A: The algorithm is similar as connectivity, except that when a cycle is detected, if it is odd, then we output "No".


### 6.3 Spanners

- Q: An $\alpha$-spanner of a graph $G=(V ; E)$ is a subgraph $H=\left(V ; E^{\prime}\right)$ such that for all $u, v, d_{G}(u, v) \leq d_{H}(u, v) \leq d_{G}(u, v)$. Can we compute a graph spanner is small space?
- A: Algorithm:

1. Let $E^{\prime}$ be initially empty.
2. On seeing $(u, v)$, set $E^{\prime}=E^{\prime} \cup(u ; v)$ if $d_{H}(u, v)>2 t-1$.

- Q: Can we get faster running time?
- A: Yes. Will talk about it next time.


## 7 K-center

Problem: Cluster a set of points $X=\left(x_{1}, x_{2}, \ldots, x_{m}\right)$ to clusters $c_{1}, c_{2}, \ldots, c_{k}$ with representatives $r_{1} \in c_{1}, r_{2} \in c_{2}, \ldots, r_{k} \in c_{k}$ to minimize $\max _{i} \min _{j} d\left(x_{i}, r_{j}\right)$.

Solution: We illustrate the algorithm in four steps.

1. First, a doubling algorithm (not necessary in the streaming setting):
(a) When we see $(k+1)$-st point in the input stream, suppose that the minimum pairwise distance is 1 , by rescaling the distances.
(b) Then at this point, we know that $\mathrm{OPT} \geq 1$, where $\mathrm{OPT}=$ cost of optimum clustering
(c) The algorithm now proceeds in phases. At the start of the ith phase, we know that $\mathrm{OPT} \geq d_{i} / 2\left(d_{1}=1\right)$.
(d) Phase i does the following:
i. Choose each point as the new representative if its distance from all other representatives $>2 d_{i}$.
ii. If we are about to choose the $(k+1)$-st representative, stop this phase, and let $d_{i+1}=2 d_{i}$.
iii. If the algorithm ends with $w<k$ clusters, we are good. and we can just add in some arbitrary cluster. This additional cluster could be empty or not.
2. The concept of summary. A summary of a stream $\sigma=\left\langle x_{1}, x_{2}, \ldots\right\rangle$ with $x_{i} \in U$ is any set $S \subseteq U$. The cost of a summary is

$$
\Delta(\sigma, S)=\max _{x \in \sigma} \min _{y \in S} d(x, y)
$$

We say $\Delta$ is a metric cost if for any streams $\sigma, \sigma \circ \pi$ and their respective summaries, $S, T$, we have:

$$
\Delta(S \circ \pi, T)-\Delta(\sigma, S) \leq \Delta(\sigma \circ \pi, T) \leq \Delta(S \circ \pi, T)+\Delta(\sigma, S)
$$

(The more precise notation for $\Delta(S \circ \pi, T)$ should be $\Delta(\sigma(S) \circ \pi, T)$ where $\sigma(S)$ is the input where every input $x \in \sigma$ is replaced by the corresponding ]hatx generated from $S$ which best represents $x$ ).
3. The concept of threshold algorithm. For a summarization cost function $\Delta$, $\Delta$ has an $\alpha$-approximate threshold algorithm $\mathcal{A}$, such that $\mathcal{A}$ takes as input an estimate $t$ and a stream $\sigma$, and
(a) Either produces a summary $S$ (where $|S|=k$ ) s.t. $\Delta(\sigma, S) \leq \alpha t$
(b) Or fail, proving that $\forall T$ (where $|T|=k$ ), $\Delta(\sigma, T)>t$

Note that the doubling algorithm is a constant-approximate threshold algorithm. Let $\alpha$ be this constant.
4. Guha's (2009) $k$-center algorithm:
(a) Keep $p$ (such that $(1+\epsilon)^{p}=\alpha / \epsilon$ ) instances of the $\mathcal{A}$ algorithm with contiguous geometrically increasing thresholds of the form $(1+\epsilon)^{i}$ running in parallel.
(b) Whenever $q \leq p$ of instances fail, start up next $q$ instances of $\mathcal{A}$ using the summaries from the failed instances the "replay" the stream. When an instance fails, kill all the other instances with a lower threshold. (Note that the threshold goes up from $(1+\epsilon)^{i} \rightarrow(1+\epsilon)^{i+p}$.
(c) At the end of input, output the summary from the lowest threshold alive instances of $\mathcal{A}$.

### 7.1 Analysis

Idea: we run multiple copies (but controlled in number) of the algorithm corresponding to different estimates of the final cost and we use a "stream-strapping" procedure to use partially completed summarization for a certain estimate to create summarization for a different estimate of cost.

Suppose the final threshold, say, $t$, has been raised by $j$ times. Let $\sigma_{i}$ and $S_{i}$ be the $i$-th portion of the stream and the summary of $\sigma_{i}$ in phase $i$, respectively. Let $T$ be the final summary. The cost of the final output $=$

$$
\begin{aligned}
\Delta\left(\sigma_{1} \circ \sigma_{2} \circ \ldots \circ \sigma_{j} \circ \pi, T\right) & \leq \Delta\left(S_{j} \circ \pi, T\right)+\Delta\left(\sigma_{1} \circ \sigma_{2} \circ \ldots, S_{j}\right) \\
& \leq \alpha t+\Delta\left(\sigma_{1} \circ \sigma_{2} \circ \ldots \circ \sigma_{j}, S_{j}\right) \\
& \leq \alpha t+\Delta\left(S_{j-1} \circ \sigma_{j}, S_{j}\right)+\Delta\left(\sigma_{1} \circ \sigma_{2} \circ \ldots \circ \sigma_{j-1}, S_{j-1}\right) \\
& \leq \alpha t+\epsilon t+\epsilon^{2} t+\ldots \\
& \leq \alpha t+(\epsilon /(1-\epsilon)) t .
\end{aligned}
$$

On the other hand, $\mathrm{OPT}>t /(1+\epsilon)$.

## 8 Edit Distances

### 8.1 Edit Distance to Monotonicity

Problem: Compute the minimum \#elements to delete such that the rest is a monotonically increasing sequence.

First talk about the method used in [GJKK07], without analysis. And then talk about the improved method used in [EJ08], with analysis.

Intuition: for each element $e$, if there are quite a lot elements in a certain range before it are inverse of it, then probably $e$ should be deleted. To further improve this, we should not not count such "reverse" caused by a (big) element that is actually been deleted (use an example in the [GJKK07] paper). That is why in [EJ08] whether $e \in R$ also depends on previous elements in $R$.

### 8.1.1 The first algorithm

Define a set $R_{\delta}$ containing all indices $i$ that are the right endpoints of an interval where more than a $\delta$-fraction of the elements lie in $\operatorname{inv}(i)$. Let $R=R_{1 / 2}$.
Thm: $|R|$ is a good estimate of ED, that is, for all $\delta \leq 1 / 2$,

$$
e d(\sigma) / 2 \leq|R| \leq e d(\sigma) / \delta
$$

Show two examples showing that this Thm is tight in terms of the estimator $R$ we choose.

### 8.1.2 The second algorithm

The algorithm. We define a new estimator $R$ :

1. $i-1 \in \operatorname{inv}(i)$, or,
2. there is an interval $I=[j, i-1]$ such that $|\operatorname{inv}(i) \cap I|>|R \cap I|$.

We say $R$ is total if $\forall i \notin R, i$ does not have a witness. We have the following theorem
Thm: $|R| \leq e d(\sigma)$. Moreover if $R$ is total then $|R| \geq \frac{1}{2} e d(\sigma)$.
Proof idea: For the left side, show that if we link each item in $R$ to its witness, the resulting graph will be a set of decreasing paths which can lower bound the ED.

For the right side, we can show that we can delete $2|R|$ elements to get a monotonic increasing sequence by deleting $[j+1, i-1]$ where $i \in R, i-1 \notin R$ and $j$ is the largest index such that $j<i$ and $j \notin R \cup i n v(i)$.

The analysis. Direct comparing $|\operatorname{inv}(i) \cap I|$ and $|R \cap I|$ is very difficult. We need a novel method to test majority: to compute the approximate median! Let $A(S, k \phi)$ be an $\epsilon$-approximation for the $\phi$-quantile of the $k$-most recent elements in stream $S$. The algorithm is as follows:

## RedTest $(j, i)$

1. Let $a=A\left(\sigma, i-j, \frac{1}{2}-\epsilon\right)$.
2. If $a \leq \sigma(i)$ return FALSE.
3. Let $a^{\prime}=A\left(R^{\prime}, i-j, \frac{1}{2}+\epsilon\right)$.
4. If $a^{\prime}=0$ then return TRUE otherwise return FALSE.

Note that we do not actually find a total $R$, but close to.
Finally, we don't have to test for all intervals $[j, i]$ but only those with length $1,2, \ldots,\left(1+\epsilon^{\prime}\right),\left(1+\epsilon^{\prime}\right)^{2}, \ldots, n$.

### 8.2 Embedding Ulam distance

Idea: turns "count", "there exists" and "majority" to $\ell_{1}, \ell_{\infty}$ and $\left(\ell_{2}\right)^{2}$. Our host space: $\bigoplus_{\left(l_{p}\right)^{p}}^{d} \bigoplus_{l_{\infty}}^{O(\log d)} l_{1}^{2 d}$ where $p=1+\epsilon$.
Definition.

1. $P_{a k}$ : the $k$ symbols that appear in the $k$ positions before $a$.
2. $\varphi_{a k}$ : the $0 / 1$ incidence vector of $P_{a k}$ scaled by $1 / 2 k$.
3. $\varphi_{a}=\bigoplus_{k \in K} \varphi_{a k}(P)$ where $K=\left\{(1+\gamma)^{i}\right\}\left(0 \leq i \leq \log _{1+\gamma} d\right)$.
4. $\varphi=\bigoplus_{a \in[d]} \varphi_{a}(P)$
$d_{p^{p}, \infty, 1}(\varphi(P), \varphi(Q))=\sum_{a \in[d]}\left(d_{\infty, 1}\left(\varphi_{a}(P), \varphi_{a}(Q)\right)\right)^{p}=\sum_{a \in[d]}\left(\max _{k \in K}\left\|\varphi_{a k}(P), \varphi_{a k}(Q)\right\|_{1}\right)^{p}$
Definition: For $P, Q \in U l a m_{d}$, and $0<\delta \leq 1 / 2$. Let $T_{\delta}$ be the set of all $a \in \Sigma$ for which there exists $k \in[d]$ such that the symmetric difference $\left|P_{a k} \Delta Q_{a k}\right|>2 \delta k$.
Lem:

$$
\frac{1}{2} e d(P, Q) \leq\left|T_{\delta}\right| \leq \frac{4}{\delta} e d(P, Q)
$$

Thm: $\varphi$ is an embedding with constant distortion.

## Proof:

1. For each $a \in T_{1 / 2}$, there exists $k \in[d]$ such that $P_{a k} \Delta Q_{a k}>2 \cdot 1 / 2 \cdot k$. Thus there exist $k^{\prime}$ such that $\left\|\varphi_{a k^{\prime}}(P), \varphi_{a k^{\prime}}(Q)\right\|_{1}=\frac{1}{2 k^{\prime}}\left|P_{a k^{\prime}} \Delta Q_{a k^{\prime}}\right|=\Omega(1)$.
2. On the other hand,

$$
\begin{aligned}
d_{p^{p}, \infty, 1}(\varphi(P), \varphi(Q))= & \sum_{a \in[d]}\left(\max _{k \in K}\left\|\varphi_{a k}(P), \varphi_{a k}(Q)\right\|_{1}\right)^{p} \\
\leq & \sum_{a \in[d]} \max _{k \in[d]}\left\|\varphi_{a k^{\prime}}(P), \varphi_{a k^{\prime}}(Q)\right\|_{1}^{1+\epsilon} \\
& (\operatorname{relax} k \in K \text { to } k \in[d]) \\
= & \sum_{a \in[d]} \max _{k \in[d]}\left[\frac{1}{2 k}\left|P_{a k} \Delta Q_{a k}\right|\right]^{1+\epsilon} \\
\leq & 1+\sum_{j=1}^{\log d}\left(2^{-j+1}\right)^{1+\epsilon} \cdot\left|T_{2^{-j}}\right|
\end{aligned}
$$

(break the contribution of $a \in \Sigma$ into buckets of the form $\left[2^{-j}, 2^{-j+1}\right]$ )
$\leq 1+\sum_{j=1}^{\log d}\left(2^{-j+1}\right)^{1+\epsilon} \cdot 2^{j+2} e d(P, Q)$
$\leq 1 / \epsilon \cdot e d(P, Q)$

## 9 Lower Bounds

### 9.1 ST-connectivity

- Q: Given an input graph stream, are two specific vertices $s$ and $t$ connected? What is the minimum space needed?
- A: $\Omega(n)$. By reduction from DISJ. Let $(x, y)$ be an instance of DISJ. Construct an instance of CONN on a graph with vertex set $\left(s, t, v_{1}, v_{2}, \ldots, v_{n-2}\right)$, and edges $A=\{(s, v): v \in x\}$ and $B=\{(v, t): v \in x\}$. It is easy to see that $\operatorname{DISJ}(x, y)=1$ iff $s, t$ are connected in the graph.


### 9.2 Perfect Matching

- Q: Given an input graph stream, does it have a perfect matching of size $n / 2$ ?
- A: $\Omega\left(n^{2}\right)$. By reduction from $\operatorname{INDEX.~Let~}(x, k)\left(x \in\{0,1\}^{n^{2}}, k \in\left[n^{2}\right]\right)$ be an instance of INDEX. Construct an instance of PM on a graph with vertex set $\cup_{i=1}^{n}\left(a_{i}, b_{i}, u_{i}, v_{i}\right)$, and edges $A=\left\{\left(u_{i}, v_{j}\right)\right.$ for each $\left.x_{f(i, j)}=1\right\}$ where $f(i, j)=(n-1) i+j$ and $B=\left\{\left(a_{l}, u_{l}\right)\right.$ for each $\left.l \neq i^{*}\right\} \cup\left\{\left(b_{l}, v_{l}\right)\right.$ for each $\left.l \neq j^{*}\right\} \cup\left\{a_{i^{*}}, b_{j^{*}}\right\}$ where $f\left(i^{*}, j^{*}\right)=k$. It is easy to see that $\operatorname{INDEX}(x, k)=1$ iff the graph has a perfect matching.


## $9.3 \quad F_{k}$

- Q: What is the space lower bounds to estimate $F_{k}$ ?
- A: $\Omega\left(n^{1-(2+\gamma) / k}\right)$ for arbitrarily small $\gamma$. By reduction from miltiparty DISJ (the promised version).
Let $\mathcal{A}$ be an $s$-space data stream algorithm that approximates $F_{k}$ within $1 \pm \epsilon$ multiplicative error with confidence $1-\delta(0<\delta<1 / 4)$. We use $\mathcal{A}$ to construct a $\delta$-error one-way protocol for DISJ $_{n, t}$, where $t=((1+3 \epsilon) n)^{1 / k}$ by partitioning the stream to $t$ parts and feeding each player one part. If there exists an intersection, then $F_{k} \geq t^{k}=(1+3 \epsilon) n$. Otherwise $F_{k} \leq n$. So a $(1+\epsilon)$ estimation of $F_{k}$ can solve DISJ $_{n, t}$. Therefore we get a protocol that solves DISJ $_{n, t}$ with communication $s(t-1)+1$. On the other hand we know that any protocol that solves DISJ $_{n, t}$ has communication at least $n / t^{1+\gamma}$. Thus we have that $s=\Omega\left(n^{1-(2+\gamma) / k}\right)$.


## 10 Some Other Topics

1. Histograms: The input is $\left\langle x_{1}, x_{2}, \ldots, x_{m}\right\rangle \in[n]^{m}$. Goal: Determine $B$ bucket histogram $H:[m] \rightarrow \mathbb{R}$ minimizing $\sum_{i \in[m]}\left(x_{i}-H(i)\right)^{2}$.

Result: $(1+\epsilon)$ estimation in $\tilde{O}\left(B^{2} \epsilon^{-1}\right)$ space. See the paper Fast, SmallSpace Algorithms for Approximate Histogram Maintenance by Gilbert et. al.
2. Transpositions: The input is $\left\langle x_{1}, x_{2}, \ldots, x_{m}\right\rangle \in[n]^{m}$. Goal is to estimate number of transpositions $\left|\left\{i<j: x_{i}>x_{j}\right\}\right|$.
Result: $(1+\epsilon)$ estimation in $\tilde{O}\left(\epsilon^{-1}\right)$ space. See the paper Counting inversions in lists by A. Gupta and F. Zane. And an improvement by G. Cormode, F. Korn, S. Muthukrishnan, and D. Srivastava. Space- and timeefficient deterministic algorithms for biased quantiles over data streams. PODS 2006.

