# CS49/249 (Randomized Algorithms), Spring 2021 : Lecture 20 

Topic: Streaming V : Counting Distinct Elements II

Disclaimer: These notes have not gone through scrutiny and in all probability contain errors.<br>Please discuss in Piazza/email errors to deeparnab@dartmouth.edu

- In the previous lecture, we saw that Flajolet-Martin gives a $O(1)$-approximation to the number of distinct elements, even with the use of pairwise independent hash functions. What if we desire a better approximation? Note that our proof of the $O(1)$ approximation was not via the usual "unbiasedestimate + low-variance + medians-of-means" method. Indeed, the estimate is not unbiased. As I mentioned in the previous lecture, Flajolet and Martin actually proved that a scaled version of their estimator is indeed (close to) an unbiased one and they also bound their variance, and also prove that taking averages tends to give a $(1 \pm \varepsilon)$ approximation to $F_{0}$. In this lecture, we will see a different modification of the algorithm which gives an $(1 \pm \varepsilon)$-estimate. This algorithm is one of three $F_{0}$ estimation algorithms in a paper by Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, D. Sivakumar, and Luca Trevisan.
- The main idea stems from an observation made last lecture. Let us recall the definitions. For any integer $1 \leq r \leq L$ (where $L:=\lceil\lg n\rceil$ ) and every element $e$ in the stream, we denote $X_{e, r}=1$ if $h(e)$ contained $\geq r$ trailing zeros, and $X_{e, r}=0$ otherwise. We denoted $Y_{r}:=\sum_{e \in D} X_{e, r}$ as the number of distinct elements in the data stream that have $\geq r$ trailing zeros. We observed that (and this needed only pairwise independence in the hash family)

$$
\begin{equation*}
\text { For all } 1 \leq r \leq L, \quad \operatorname{Exp}\left[Y_{r}\right]=\frac{d}{2^{r}} \quad \text { and } \quad \operatorname{Var}\left[Y_{r}\right]=\frac{d}{2^{r}} \cdot\left(1-\frac{1}{2^{r}}\right) \tag{1}
\end{equation*}
$$

And thus, the scaled random variable $Z_{r}:=2^{r} Y_{r}$ is an unbiased estimate of $d$. Therefore, by Chebyshev, if we want an $(1 \pm \varepsilon)$-estimate to $d$, we need to focus on the $r$ such that $2^{r} \leq \varepsilon^{2} d$. More precisely, Chebyshev gives us

$$
\begin{equation*}
\text { For all } 1 \leq r \leq L, \operatorname{Pr}\left[\left|2^{r} Y_{r}-d\right| \geq \varepsilon d\right]=\operatorname{Pr}\left[\left|Y_{r}-\frac{d}{2^{r}}\right| \geq \frac{\varepsilon d}{2^{r}}\right]<\frac{2^{r}}{\varepsilon^{2} d} \tag{2}
\end{equation*}
$$

Why not then just pick an $r$ which is small enough? The flip is the size required to evaluate $Y_{r}$. How do we keep track of $Y_{r}$ ? For instance, if $r=1$, then $Y_{r}$ is the number of distinct elements which have $\geq 1$ leading zeros. But that's going to be $\approx \frac{1}{2}$ the elements. And note that we need to actually store the elements to keep track : when the same copy of a previously seen element arrives, we need to make sure not to double count. So, we need to make sure $r$ is "big enough" such that $Y_{r}$ itself is not too large. In particular, we want to choose $r$ such that $2^{r} \approx \varepsilon^{2} d$; but this can't be a priori fixed. One needs to be more ingenious.

Remark: Before moving on, however, note that a "two pass" algorithm leaps out at us. That is, suppose we could make two scans across the whole stream. Then, in the first scan/stream find the estimate $r$ by Flajolet-Martin with the guarantee $\frac{d}{8} \leq 2^{r} \leq 4 d$. In the second pass, keep track of $Y_{r}$. That is, explicitly store all elements with $\geq r$ trailing zeros. We expect this number to be $O(1)$, and so the total storage will be $O(\log n)$ bits. The variance-by-squaredmean is also a constant, which means $O\left(\frac{1}{\varepsilon^{2}} \ln (1 / \delta)\right)$ parallel repetitions will suffice to give us an

## $(\varepsilon, \delta)$-approximation to $d$.

- The idea behind the one-pass algorithm is to start with a "small" $r$ and then keep bumping it up as we go. Since we are only in the insertion-only setting, the number of distinct elements cannot go down. To explain this, let us go back to our "wasteful" description of the Flajolet-Martin algorithm. Imagine, we have $L$ counters, or rather more precisely, buckets. When an item $e$ arrives, we evaluate $r(e)$ and then plonk it into the $C[r(e)]$ th bucket if it is not already present. The presence check can be done via a linear scan (at the end, these buckets will be small, so no real need to be clever here, but one could use a binary search tree or a hash table). Note that for any $r$, the sum of the sizes of the buckets with index bigger than equal to $r$ is precisely $Y_{r}$. Worth writing this explicitly:

$$
\text { For any } 1 \leq r \leq L, \quad Y_{r}=\sum_{j \geq r}|C[j]|
$$

As described so far, the algorithm stores all the elements. But here is the kicker: at any point of time, the algorithm only stores buckets with index $\geq r$ such that $Y_{r} \leq \frac{c}{\varepsilon^{2}}$ for some constant $c$. In doing so, it ensures that the algorithm never stores more than $\frac{c}{\varepsilon^{2}}$ elements, and thus never uses more than $O\left(\frac{c}{\varepsilon^{2}}\right)$-words of memory (or $O\left(\frac{c \log n}{\varepsilon^{2}}\right)$ bits).
More precisely, the algorithm maintains a bucket index rmin initialized to 1 and only stores the buckets $C[j]$ with $j \geq \mathrm{rmin}$. So, when an item $e$ arrives and $r(e)<\mathrm{rmin}$, the algorithm ignores the element. The algorithm also maintains $Y_{r \min }$ as defined above. Initially, this is 0 . If $r(e) \geq r m i n$, then it adds $e$ to $C[r(e)]$ if $e$ is not already present. This increases $Y_{\mathrm{rmin}}$ by 1 . If $Y_{\mathrm{rmin}}$ exceeds $\frac{c}{\varepsilon^{2}}$, then the algorithm deletes $C[r m i n]$ from memory, updates $Y_{\mathrm{rmin}}$, and increments rmin by 1 . At the end of the stream, this $r$ min is the " $r$ " we wanted: the algorithm outputs est $\leftarrow 2^{r \text { min }} \cdot Y_{\text {rmin }}$.

```
procedure BASIC-BJKST:
    Choose \(h:[n] \rightarrow\left[2^{L}\right]\) from a strongly universal hash family as in FLAJOLET-MARTIN.
    \(C[1: L]\) is a list of buckets.
    \(\mathrm{rmin} \leftarrow 1, Y_{\mathrm{rmin}} \leftarrow 0\).
    for arrival of element \(e\) do:
        Evaluate \(r(e)\) : the number of trailing zeros in \(h(e)\).
        if \(r(e)<\mathbf{r m i n}\) then:
            Ignore this element \(e\).
        else:
            If \(e\) is already present in \(C[r(e)]\) ignore this element.
            Otherwise, add \(e\) to \(C[r(e)]\) and update \(Y_{\mathrm{rmin}} \leftarrow Y_{\mathrm{rmin}}+1 . \triangleright\) Note that \(Y_{\mathrm{rmin}}=\)
\(\sum_{j \geq \mathrm{rmin}}|C[j]|\) is maintained.
            \(\triangleright\) This increase in \(Y_{\mathrm{rmin}}\) may make it \(>\frac{c}{\varepsilon^{2}}\). The next while loop fixes this.
            while \(Y_{\text {rmin }}>\frac{c}{\varepsilon^{2}}\) do:
                Delete \(C[\mathrm{rmin}]\) and update \(Y_{\mathrm{rmin}} \leftarrow Y_{\mathrm{rmin}}-|C[\mathrm{rmin}]|\).
                    Update \(\mathrm{rmin} \leftarrow \mathrm{rmin}+1 . \triangleright\) Note that \(Y_{\mathrm{rmin}}=\sum_{j \geq \mathrm{rmin}}|C[j]|\) is maintained.
            \(\triangleright\) Note: if at the beginning of the while loop, \(C[\mathrm{rmin}]>0\), then it runs only
once. But Line 16 may lead rmin to point to a bucket with \(C[\mathrm{rmin}]=0\).
    return est \(\leftarrow 2^{\text {rmin }} \cdot Y_{\text {rmin }}\).
```

Observation 1. Since items are only inserted, throughout the algorithm $Y_{\text {rmin }}$ equals $\sum_{e \in D} X_{e, r m i n}$, which is the number of elements in $D$ whose hash has $\geq$ rmin trailing zeros.

## - Analysis of Quality.

Theorem 1. The estimate est returned by BASIC-BJKST satisfies $(1-\varepsilon) d \leq$ est $\leq(1+\varepsilon) d$ with probability $\geq \frac{3}{4}$.

Proof. If you have followed the intuition of the algorithm, then you probably see that by design the worry of rmin being "too small" is allayed. The algorithm never stores more than $O\left(1 / \varepsilon^{2}\right)$ items. The worry, if any, is whether rmin became too big. Or rather, can $2^{r m i n}>\varepsilon^{2} d$ ? The answer is no. Consider the time when rmin is incremented from some $k$ to $k+1$. At that point, we must have $Y_{k}>\frac{c}{\varepsilon^{2}}$. Can $k$ be such that $2^{k}>\varepsilon^{2} d ?$ No, because $\operatorname{Exp}\left[Y_{k}\right]=\frac{d}{2^{k}}$ with variance also of that order, and so whp $2^{k} \approx \frac{\varepsilon^{2} d}{c}$ in this case. So, the final rmin we output, will almost sure be such that $\operatorname{Var}\left[Y_{\mathrm{rmin}}\right] / \operatorname{Exp}^{2}\left[Y_{\mathrm{rmin}}\right]$ is small, and therefore, the estimate $2^{r m i n} \cdot Y_{r m i n}$ should be a good one.
To make the above proof formal, there is a bit of care needed. Note that we can't simply argue about $Y_{\text {rmin }}$ as it is a random variable indicated by a random variable. While, (1) is for random variables with a fixed index. So, a bit more care is needed. To do so, we go over all $L$ possibilities of rmin, and argue that the probability of something bad is happening is small. What is bad? Well, the bad event for us is the following: $\mathrm{rmin}=k$ but $2^{k} Y_{k}$ is not a good estimate. Let's define this as an event:

$$
\text { For } 1 \leq k \leq L, \quad \mathcal{E}_{k}:=\left\{\operatorname{rmin}=k \wedge\left|2^{k} Y_{k}-d\right| \geq \varepsilon d\right\}
$$

We want to prove, $\operatorname{Pr}\left[\bigvee_{k=1}^{L} \mathcal{E}_{k}\right]<\frac{1}{4}$. We will proceed by union bound. We break the $k$ 's into "small" and "large". To define this, let $k_{*}$ to be the largest integer such that $\frac{2^{k_{*}}}{d} \leq \frac{\varepsilon^{2}}{16}$. Then, using Equation (2) and the union bound, we get

$$
\begin{equation*}
\operatorname{Pr}\left[\bigvee_{k \leq k_{*}} \mathcal{E}_{k}\right] \underbrace{\leq}_{\operatorname{Pr}[A \wedge B] \leq \operatorname{Pr}[B]} \operatorname{Pr}\left[\bigvee_{k \leq k_{*}}\left\{\left|2^{k} Y_{k}-d\right| \geq \varepsilon d\right\}\right] \underbrace{\leq}_{\text {Union Bound and (2) }} \frac{1}{\varepsilon^{2} d} \sum_{k=1}^{k_{*}} 2^{k} \leq \frac{1}{8} \tag{3}
\end{equation*}
$$

What about the "large" $k$ 's? Well, we bound $\operatorname{Pr}\left[\bigvee_{k>k_{*}} \mathcal{E}_{k}\right] \leq \operatorname{Pr}\left[\mathrm{rmin}>k_{*}\right]$. Which means that $Y_{k_{*}}>\frac{c}{\varepsilon^{2}}$. By our choice of $k_{*}$, we know that $\frac{2^{k_{*}}}{d}>\frac{\varepsilon^{2}}{32}$. That is, $\operatorname{Exp}\left[Y_{k^{*}}\right]=\frac{d}{2^{k^{*}}}<\frac{32}{\varepsilon^{2}}$, and thus just Markov gives us that

$$
\operatorname{Pr}\left[Y_{k_{*}}>\frac{c}{\varepsilon^{2}}\right] \leq \frac{\operatorname{Exp}\left[Y_{k^{*}}\right]}{c / \varepsilon^{2}}<\frac{32}{c}<\frac{1}{8} \text { if } c \text { is large enough. }
$$

Putting everything together, we get

$$
\begin{equation*}
\operatorname{Pr}\left[\bigvee_{k>k_{*}} \mathcal{E}_{k}\right] \underbrace{\leq}_{\operatorname{Pr}[A \wedge B] \leq \operatorname{Pr}[B]} \operatorname{Pr}\left[\mathrm{rmin}>k_{*}\right]=\operatorname{Pr}\left[Y_{k^{*}}>\frac{c}{\varepsilon^{2}}\right] \underbrace{\leq}_{\text {if } c \text { large enough }} \frac{1}{8} \tag{4}
\end{equation*}
$$

Union bounding over Equation (3) and Equation (4) gives that $\operatorname{Pr}\left[\bigvee_{k=1}^{L} \mathcal{E}_{k}\right]<\frac{1}{4}$.

- Time and Space: a space saving trick by BJKST. Per update, the algorithm spends time evaluating $r(e)$. After that, the majority of the time taken is in checking if $e$ is already in $C[r(e)]$. Since $|C[r(e)]| \leq Y_{\mathrm{rmin}} \leq \frac{c}{\varepsilon^{2}}$, this takes at most $O\left(\frac{1}{\varepsilon^{2}}\right)$ time. One can do much better though : either $O\left(\lg \frac{1}{\varepsilon}\right)$ by storing $C[r(e)]$ as a binary search tree, or even $O(1)$ amortized time by just hashing. This forms the lion's share of the update time.
How about space? There is some space required to store the hash functions. This is $O(\log n)$ bits or $O(1)$ words. The bigger usage of space is the buckets. Note, we only store the buckets for $j \geq$ rmin, and we maintain $\left|Y_{\text {rmin }}\right| \leq \frac{c}{\varepsilon^{2}}$. Therefore, the space required is $O\left(\frac{\log n}{\varepsilon^{2}}\right)$. This doesn't sound too bad, till you compare with the space required by FLAJOLET-MARTIN: ignoring the hash function, the space usage of that algorithm was only $\lg \lg n+O(1)$ bits. Next, we discuss a space-saving trick by BJKST which takes motivation from the birthday paradox.
The main idea is again hashing. Note that we don't really need to store the element $e$ in the bucket. We just need to make sure when another copy of $e$ arrives we don't count it again. So instead of storing $e$, we just store a hash $g(e)$ for some hash function $g:[n] \rightarrow[s]$. The question is how big does $s$ need to be? If we wanted no collisions at all, that is we wanted $g$ to be perfect, then the constructions we studied were of size $s=O(n)$. But this defeats the purpose - storing $g(e)$ would take the same amount of space.
But then we realize that the elements we desire no collision on are the ones ever present in the buckets. So we ask ourselves : how many distinct elements $e \in[n]$ are ever present in the buckets? Crudely, for every $r$, the size of the bucket $C[r]$ is $\leq \frac{c}{\varepsilon^{2}}$, and there are $L=O(\lg n)$ possible such $r$ 's. Thus, in the run of BASIC-BJKST, if we consider the (random) set $S \subseteq[n]$ that is hashed into the buckets, this size $|S|=O\left(\frac{\lg n}{\varepsilon^{2}}\right)$. Therefore, we don't need the range of $g$ to be large. If $g:[n] \rightarrow[s]$ where $s=\frac{b \lg ^{2} n}{\varepsilon^{4}}$, then by a birthday-paradox style argument, the probability there is a collision among the elements is $S$ is $\leq \frac{1}{12}$ for a large enough $b$. One can add this to the failure probability, and get the success probability of the full algorithm to be $\geq \frac{2}{3}$.

```
procedure BJKST:
    Choose \(h:[n] \rightarrow\left[2^{L}\right]\) from a strongly universal hash family as in Flajolet-Martin.
    Choose \(g:[n] \rightarrow\left[\frac{b \log ^{2} n}{\varepsilon^{4}}\right]\) from a UHF.
    \(C[1: L]\) is a list of buckets. \(\triangleright\) In reality, one uses a data structure which can dynamically
store buckets indexed via a key
    \(\mathrm{rmin} \leftarrow 1, Y_{\mathrm{rmin}} \leftarrow 0\).
    for arrival of element \(e\) do:
        Evaluate \(r(e)\) : the number of trailing zeros in \(h(e)\).
        if \(r(e)<r m i n\) then:
            Ignore this element \(e\).
        else:
            If \(g(e)\) is already present in \(C[r(e)]\) ignore this element. \(\triangleright\) Use a binary search
    tree, or another hash table
            Otherwise, add \(g(e)\) to \(C[r(e)]\) and update \(Y_{\mathrm{rmin}} \leftarrow Y_{\mathrm{r} \min }+1 . \triangleright\) Note that \(Y_{\mathrm{rmin}}=\)
    \(\sum_{j \geq r \min }|C[j]|\) is maintained.
            \(\triangleright\) This increase in \(Y_{\mathrm{rmin}}\) may make it \(>\frac{c}{\varepsilon^{2}}\). The next while loop fixes this.
            while \(Y_{\mathrm{rmin}}>\frac{c}{\varepsilon^{2}}\) do:
            Delete \(C[r \mathrm{rmin}]\) and update \(Y_{\mathrm{rmin}} \leftarrow Y_{\mathrm{rmin}}-|C[r m i n]|\).
            Update \(\mathrm{rmin} \leftarrow \mathrm{rmin}+1 . \triangleright\) Note that \(Y_{\mathrm{rmin}}=\sum_{j \geq \mathrm{rmin}}|C[j]|\) is maintained.
            \(\triangleright\) Note: if at the beginning of the while loop, \(C[\mathrm{rmin}]>0\), then it runs only
    once. But Line 16 may lead rmin to point to a bucket with \(C[\mathrm{rmin}]=0\).
    return est \(\leftarrow 2^{\text {rmin }} \cdot Y_{\text {rmin }}\).
```

Lemma 1. For any subset $S \subseteq[n]$, the probability there exists $e, e^{\prime} \in S$ such that $g(e)=g\left(e^{\prime}\right)$ is at most $\frac{|S|^{2} \varepsilon^{4}}{2 b \lg ^{2} n}$.

Proof. The probability of collision for a fixed pair is $\leq \frac{\varepsilon^{4}}{b \lg ^{2} n}$ since $g$ is drawn from a UHF. The lemma follows by a union bound over the $\leq|S|^{2} / 2$ different pairs.

