

# Data Streams as Random Permutations: the Distinct Element Problem


Conrado Martínez,  
Univ. Politècnica de Catalunya

VIII JMDA, Almería, Julio 2012

Joint work with:



A. Helmi



J. Lumbroso



A. Viola

# Introduction



- A **data stream** is a (very long) sequence

$$\mathcal{S} = s_1, s_2, s_3, \dots, s_N$$

of items  $s_i$  drawn from some (large) domain  $\mathcal{U}$ ,  $s_i \in \mathcal{U}$

- The goal: to compute  $\theta = \theta(\mathcal{S})$ , but there are **limitations** to our computational power:
  - a single pass over the sequence
  - very short time for computation on each item
  - very small auxiliary memory:  $M \ll N$ ; ideally  $M = \Theta(1)$  or  $M = \mathcal{O}(\log N)$
  - no statistical hypothesis on the data

# Introduction



There are lots of applications for this **data stream model**:

- Network traffic analysis  $\Rightarrow$  DoS/DDoS attacks, worms, . . .
- Database query optimization
- Information retrieval  $\Rightarrow$  similarity index
- Data mining
- And many more . . .

# Introduction



We will often see  $\mathcal{S}$  as a multiset

$\{y_1^{f_1}, \dots, y_n^{f_n}\}$ ,  $f_i$  = frequency of the  $i$ th distinct element  $y_i$

Some typical problems:

- The cardinality of  $\mathcal{S}$ :  $\text{card}(\mathcal{S}) = n \leq N = |\mathcal{S}| \Leftarrow$  **This paper**
- The elements  $y_i$  such that:
  - $f_i \geq k$  (**k-elephants**)
  - $f_i \geq c \cdot N$ ,  $0 < c < 1$  (**c-icebergs**)
  - $f_i < k$  (**k-mice**)

# Introduction

Small auxiliary memory  $\Rightarrow$

Exact solution too costly (or impossible)  $\Rightarrow$

Randomized algorithms  $\Rightarrow$

Estimation  $\hat{\theta}$  of the quantity  $\theta$

- The estimator  $\hat{\theta}$  must be **unbiased**

$$E [\hat{\theta}] = \theta$$

- The estimator must be accurate (small **standard error**)

$$SE [\hat{\theta}] := \frac{\sqrt{\text{Var} [\hat{\theta}]}}{E [\hat{\theta}]} < \epsilon,$$

e.g.,  $\epsilon = 0.01$  (1%)

# Estimating the cardinality

The first ingredient:

- Map each item  $s_i$  of the stream to a value in  $(0, 1)$  using a **hash** function  $h : \mathcal{U} \rightarrow (0, 1) \Rightarrow$  **reproducible randomness**
- The multiset  $\mathcal{S}$  is mapped to a multiset

$$\mathcal{S}' = h(\mathcal{S}) = \{x_1^{f_1}, \dots, x_n^{f_n}\},$$

with  $x_i = \text{hash}(y_i)$ ,  $f_i = \#$  of  $x_i$ 's

- The set of **distinct** elements  $X = \{x_1, \dots, x_n\}$  is a set of  $n$  independent and uniformly distributed real numbers in  $(0, 1)$

## Estimating the cardinality

The second ingredient:

- Define some easily computable **observable**  $R$  which is insensitive to repetitions, that is, it only depends on the underlying set of distinct elements:

$$R = R(\mathcal{S}) = R(X)$$

- Perform the probabilistic analysis of  $R$  for a set  $X$  of  $n$  random real numbers. If

$$E_n [R] = \varphi(n)$$

then it is reasonable to assume that the expected value of  $\varphi^{-1}(R)$  will be close to  $n$ ; we will need some **correcting factor**  $\alpha$  to get an (asymptotically) unbiased estimator

$$E_n \left[ \alpha \varphi^{-1}(R) \right] = n + \text{l.o.t.}$$

## Probabilistic Counting

- For instance, in Flajolet & Martin's Probabilistic Counting (1985) the observable  $R$  is the length of the longest prefix  $0.0^{R-1}1$  such that all prefixes  $0.0^k1$  appear among the hashed values, for  $0 \leq k \leq R - 1$
- $R$  is easy to compute and it does not depend on repetitions

$$E_n [R] \approx \log_2 n$$

and

$$E_n [\alpha 2^R] = n + o(n)$$

for

$$\alpha^{-1} = \frac{e^\gamma \sqrt{2}}{3} \prod_{k \geq 1} \left( \frac{(4k+1)(2k+1)}{2k(4k+3)} \right)^{(-1)^{v(k)}} \approx 0.77351 \dots$$




## Other estimators

- Based on **bit patterns**
  - **LogLog**: Durand, Flajolet (2003)
  - **HyperLogLog**: Flajolet, Fusy, Gandouet, Meunier (2007)
- Based on **order statistics** (e.g., the  $k$ th smallest in the set of distinct hash values)
  - Bar-Yossef, Kumar & Sivakumar (2002)
  - Bar-Yossef, Jayram, Kumar, Sivakumar & Trevisan (2002)
  - Giroire (2005, 2009)
  - Chassaing & G erin (2006)
  - Lumbroso (2010)

## Our estimator: Recordinality

- RECORDINALITY counts the number of records (local maxima) in the sequence



- It depends in the underlying **permutation** of the first occurrences of distinct values, very different from the other estimators

## Our estimator: Recordinality

- $\sigma(i)$  is a **record** of the permutation  $\sigma$  if  $\sigma(i) > \sigma(j)$  for all  $j < i$
- It is well known that the number  $r$  of records satisfies

$$E_n [r] = \log n + \mathcal{O}(1)$$

hence we anticipate that  $e^r$  should give us an estimate of  $n$

- The notion is generalized to **k-records**:  $\sigma(i)$  is a  $k$ -record if  $\sigma(i)$  is among the  $k$  largest elements in  $\sigma(1), \dots, \sigma(i)$

## Our estimator: Recordinality

```
procedure RECORDINALITY( $\mathcal{S}$ ,  $k$ )  
  fill  $T$  with the first  $k$  distinct elements (hash values)  
  of the stream  $\mathcal{S}$   
   $r \leftarrow k$   
  for all  $s \in \mathcal{S}$  do  
     $x \leftarrow \text{hash}(s)$   
    if  $x > \min(T) \wedge x \notin T$  then  
       $r \leftarrow r + 1$ ;  $T \leftarrow T \cup \{x\} \setminus \min(T)$   
    end if  
  end for  
  return  $Z = \alpha_k \cdot e^r$   $\triangleright \alpha_k$  is a correcting factor  
end procedure
```

Memory:  $k$  hash values ( $k \log n$  bits) + 1 counter ( $\log \log n$  bits)

## Our estimator: Recordinality

### Theorem (Helmi, Martínez and Panholzer, 2012)

Let  $r_k$  denote the number of  $k$ -records in a permutation of size  $n$ . The exact distribution of  $r_k$  is

$$\text{Prob}_n \{r_k = j\} = \begin{cases} \llbracket n = j \rrbracket & \text{if } k > n, \\ k^{j-k} \frac{k!}{n!} \begin{bmatrix} n - k + 1 \\ j - k + 1 \end{bmatrix} & \text{if } k \leq j \leq n \end{cases}$$

$\begin{bmatrix} n \\ j \end{bmatrix}$  = signless Stirling numbers of the first kind;  $\llbracket P \rrbracket = 1$  if  $P$  true, = 0 otherwise

## Our estimator: Recordinality

- The expected value of  $r_k$  is  $k \log(n/k) + \text{l.o.t.}$ ; it is reasonable then to assume that for

$$Z := k \exp(\alpha_k \cdot r_k)$$

we should have  $E_n [Z] \sim n$  for some suitable correcting factor  $\alpha_k$

- We can use the formula for  $\text{Prob}_n \{r_k = j\}$  to explicitly compute  $E_n [Z]$  and to determine  $\phi$ , and then compute the standard error

## Our estimator: Recordinality

### Theorem

*The RECORDINALITY estimator*

$$Z := k \left( 1 + \frac{1}{k} \right)^{r_k - k + 1} - 1$$

*is an **unbiased** estimator of  $n$ :  $E_n [Z] = n$ .*

## Our estimator: Recordinality

### Theorem

*The accuracy of RECORDINALITY, expressed in terms of standard error, asymptotically satisfies*


$$SE_n [Z] \sim \sqrt{\left(\frac{n}{ke}\right)^{\frac{1}{k}} - 1}$$




## Our estimator: Recordinality

For practical values of  $n$ , even for small  $k$ , the estimates may be significantly concentrated.

For instance, for  $k = 10$ , the estimates are within  $\sigma$ ,  $2\sigma$ ,  $3\sigma$  of the exact count in respectively 91%, 96% and 99% of all cases.




$k = 64$




$k = 256$

500 estimates of cardinality in Shakespeare's *A Midsummer Night's Dream*; top and bottom lines (5%), centermost lines (70%); gray area (1 standard deviation)

## Other issues



Original texts



Randomly permuted texts

- RECORDINALITY does not depend on the hash values, only the relative ordering  $\Rightarrow$  **we can avoid using the hash function**, provided the distinct elements appear (for the first time) in random order
- We can combine RECORDINALITY with any of the existing estimators since they are **independent**; a suitably weighted sum of the estimations will have less variance  $\Rightarrow$  **better accuracy**

## Other issues

- The table of  $k$  largest hash values gives us a random sample of  $k$  distinct elements out of the  $n \Rightarrow$  **distinct sampling** for free
- Indeed, distinct elements “enter” the table or not according to their hash value, a random uniform number
- An easy modification allows us to have a random sample of distinct elements with expected size  $k \log(n/k) \Rightarrow$  **variable-size sampling**

## Concluding remarks

- First (?) application of combinatorics of random permutations to data stream algorithms
- Simple and elegant algorithms
- Nice combinatorics and mathematical analysis
- Many extensions to explore: sampling, sliding windows, similarity index, .....



Thanks a lot for your  
attention!

¡Gracias por vuestra  
atención!