# Mining Data Streams I 

Alexander Schönhuth

## Bielefeld University May 27, 2021

## TODAY

Overview

- MapReduce II: Recap
- Reducer Size / Replication Rate
- Graph Model / Mapping Schema
- Lower Bounds on Replication Rate
- Mining Data Streams I
- Intro: A Data Stream Management Model
- Sampling Data in a Stream
- Filtering Streams: Bloom Filters
- Counting Distinct Elements: Flajolet-Martin algorithm

Learning Goals: Understand these topics and get familiarized

## Complexity Theory for MapReduce

## MapReduce: Complexity Theory

## Idea

- Reminder: A "reducer" is the execution of a Reduce task on a single key and the associated value list
- Important considerations:
- Keep communication cost low
- Keep wall-clock time low
- Execute each reducer in main memory
- Intuition:
- The less communication, the less parallelism, so
- the more wall-clock time
- the more main memory needed
- Goal: Develop encompassing complexity theory


## Reducer Size: Informal Explanation



Reducer size: maximum length of list [ $\mathrm{v}, \mathrm{w}, \ldots \mathrm{]}$ ] after grouping keys Adopted from mmds.org

## Reducer Size

Definition [Reducer Size]:
The reducer size $q$ is the upper bound on the number of values to appear in the list of a single key.

Motivation

- Small reducer size forces to have many reducers
- Further creating many Reduce tasks implies high parallelism, hence small wall-clock time
- Sufficiently small reducer size allows to have all data in main memory


## Replication Rate

Definition [Replication Rate]:
The replication rate $r$ is the number of all key-value pairs generated by Map tasks, divided by the number of inputs.

Motivating Example

- One-pass matrix multiplication algorithm:
- Matrices involved are $n \times n$
- Reminder: Key-value pairs for $M N$ are $\left((i, k),\left(M, j, m_{i j}\right)\right), j=1, \ldots, n$ and $\left((i, k),\left(N, j, n_{j k}\right)\right), j=1, \ldots, n$
- Replication rate $r$ is equal to $n$ :
- Inputs are all $m_{i j}$ and $n_{j k}$
- For each $m_{i j}$, one generates key-value pairs for $(i, k), k=1, \ldots, n$
- For each $n_{j k}$, one generates key-value pairs for $(i, k), i=1, \ldots, n$
- Reducer size is $2 n$ : for each key $(i, k)$ there are $n$ values from each $m_{i j}$ and $n$ values from each $n_{j k}$


## Example: Similarity Join

## Situation

- Given large set $X$ of elements
- Given similarity measure $s(x, y)$ for measuring similarity between $x, y \in X$
- Measure is symmetric: $s(x, y)=s(y, x)$
- Output of the algorithm: all pairs $x, y$ where $s(x, y) \geq t$ for threshold $t$
- Exemplary input: 1 million images $\left(i, P_{i}\right)$ where
- $i$ is ID of image
- $P_{i}$ is picture itself
- Each picture is 1 MB


## Example: Similarity Join

## MapReduce: Bad Idea

- Use keys $(i, j)$ for pair of pictures $\left(i, P_{i}\right),\left(j, P_{j}\right)$
- Map: generates $\left((i, j),\left[P_{i}, P_{j}\right]\right)$ as input for
- Reduce: computes $s\left(P_{i}, P_{j}\right)$ and decides whether $s\left(P_{i}, P_{j}\right) \geq t$
- Reducer size $q$ is small: 2 MB ; expected to fit in main memory
- However, each picture makes part of 999999 key-value pairs, so

$$
r=999999
$$

- Hence, number of bytes communicated from Map to Reduce is

$$
10^{6} \times 999999 \times 10^{6}=10^{18}
$$

that is one exabyte

$$
0
$$

## Example: Similarity Join

## MapReduce: Better Idea

- Group images into $g$ groups, each of $10^{6} / g$ pictures
- Map: For each $\left(i, P_{i}\right)$ generate $g-1$ key-value pairs
- Let $u$ be group of $P_{i}$
- Let $v$ be one of the other groups
- Keys are sets $\{u, v\}$ (set, so no order!)
- Value is $\left(i, P_{i}\right)$
- Overall: $\left(\{u, v\},\left(i, P_{i}\right)\right)$ as key-value pair
- Reduce: Consider key $\{u, v\}$
- Associated value list has $2 \times \frac{10^{6}}{g}$ values
- Consider $\left(i, P_{i}\right)$ and $\left(j, P_{j}\right)$ when $i, j$ are from different groups
- Compute $s\left(P_{i}, P_{j}\right)$
- Compute $s\left(P_{i}, P_{j}\right)$ for $P_{i}, P_{j}$ from same group on processing keys $\{u, u+1\}$


## Example: Similarity Join

## MapReduce: Better Idea

- Replication rate is $g-1$
- Each input element $\left(i, P_{i}\right)$ is turned into $g-1$ key-value pairs
- Reducer size is $2 \times \frac{10^{6}}{g}$
- Number of values on list for reducer
- This yields $2 \times \frac{10^{6}}{g} \times 10^{6}$ bytes stored at Reducer node


## Example: Similarity Join

## MapReduce: Better Idea

- Example $g=1000$ :
- Input is 2 GB , fits into main memory
- Communication cost:

$$
\begin{equation*}
\underbrace{\left(10^{3} \times 999\right)}_{\text {number of reducers }} \times \underbrace{\left(2 \times 10^{3} \times 10^{6}\right)}_{\text {reducer size }} \approx 10^{15} \tag{1}
\end{equation*}
$$

- 1000 times less than brute-force
- Half a million reducers: maximum parallelism at Reduce nodes
- Computation cost is independent of $g$
- Always all-vs-all comparison of pictures
- Computing $s\left(P_{i}, P_{j}\right)$ for all $i, j$


## MapReduce: Graph Model

Goal: Proving lower bounds on replication rate as function of reducer size, for many problems. Therefore:

## Graph Model:

- Graph describes how outputs depend on inputs
- Reducers operate independently: each output has one reducer that receives all input required to compute output
- Model foundation:
- There is a set of inputs
- There is a set of outputs
- Outputs depends on inputs: many-to-many relationship


## MapReduce: Graph Model Example



Graph for similarity join with four pictures
Adopted from mmds.org

## MapReduce: Graph Model Matrix MUltiplication

Graph Model Matrix Multiplication

- Multiplying $n \times n$ matrices $M$ and $N$ makes
- $2 n^{2}$ inputs $m_{i j}, n_{j k}, 1 \leq i, j, k \leq n$
- $n^{2}$ outputs $p_{i k}:=(M N)_{i k}, 1 \leq i, k \leq n$
- Each output $p_{i k}$ needs $2 n$ inputs $m_{i 1}, m_{i 2}, \ldots, m_{i n}$ and $n_{1 k}, n_{2 k}, \ldots, n_{n k}$
- Each input relates to $n$ outputs: e.g. $m_{i j}$ to $p_{i 1}, p_{i 2}, \ldots, p_{i n}$


## MapReduce: Graph Model Matrix Multiplication II



$$
\left[\begin{array}{ll}
a & b \\
c & d
\end{array}\right]\left[\begin{array}{ll}
e & f \\
g & h
\end{array}\right]=\left[\begin{array}{ll}
i & j \\
k & l
\end{array}\right]
$$

Input-output relationship graph for multiplying $2 \times 2$ matrices
Adopted from mmds.org

## MapReduce: Mapping Schemas

A mapping schema with a given reducer size $q$ is an assignment of inputs to reducers such that

- No reducer receives more than $q$ inputs
- For every output, there is a reducer that receives all inputs required to generate the output

Consideration: The existence of a mapping schema for a given $q$ distinguishes problems that can be solved in a single MapReduce job from those that cannot.

## MApping Schema: Example

Consider computing similarity of $p$ pictures, divided into $g$ groups.

- Number of outputs: $\binom{p}{2}=\frac{p(p-1)}{2} \approx \frac{p^{2}}{2}$
- Reducer receives $2 p / g$ inputs necessary reducer size is $q=2 p / g$
- Replication rate is $r=g-1 \approx g$ :

$$
r=2 p / q
$$

$r$ inversely proportional to $q$ which is common

- In a mapping schema for reducer size $q$ :
- Each reducer is assigned exactly $2 p / g$ inputs
- In all cases, every output is covered by some reducer


## Mapping Schemas: Not all Inputs Present

Example: Natural Join $R(A, B) \bowtie S(B, C)$, where many possible tuples $R(a, b), S(b, c)$ are missing.

- Theoretically $q=|A| \cdot|C|$ (keys were $b \in B$ )
- But in practice many tuples $(a, b),(b, c)$ are missing for each $b$, so $q$ possibly much smaller than $|A| \cdot|C|$

Main Consideration: One can increase $q$ because of the missing inputs, without that inputs do no longer fit into main memory in practice

## Mapping Schemas: LOWER BOUNDS ON Replication Rate

Technique for proving lower bounds on replication rates

1. Prove upper bound $g(q)$ on how many outputs a reducer with $q$ inputs can cover
This may be difficult in some cases
2. Determine total number of outputs $O$
3. Let there be $k$ reducers with $q_{i}<q, i=1, \ldots, k$ inputs observe that $\sum_{i=1}^{k} g\left(q_{i}\right)$ needs to be no less than $O$
4. Manipulate the inequality $\sum_{i=1}^{k} g\left(q_{i}\right) \geq O$ to get a lower bound on $\sum_{i=1}^{k} q_{i}$
5. Dividing the lower bound on $\sum_{i=1}^{k} q_{i}$ by number of inputs is lower bound on replication rate

## Lower Bounds: Example All-Pairs Problem

- Recall that $r \leq 2 p / q$ was upper bound on replication rate for all-pairs problem
- Here: Lower bound on $r$ that is half the upper bound


## Lower Bounds: Example All-Pairs Problem

- Steps from slide before:
- Step 1: reducer with $q$ inputs cannot cover more than $\binom{q}{2} \approx q^{2} / 2$ outputs
- Step 2: overall $\binom{p}{2} \approx p^{2} / 2$ outputs must be covered
- Step 3: So, the inequality approximately evaluates as

$$
\sum_{i=1}^{k} q_{i}^{2} / 2 \geq p^{2} / 2 \quad \Longleftrightarrow \quad \sum_{i=1}^{k} q_{i}^{2} \geq p^{2}
$$

- Step 4: From $q \geq q_{i}$, we obtain

$$
q \sum_{i=1}^{k} q_{i} \geq p^{2} \quad \Longleftrightarrow \quad \sum_{i=1}^{k} q_{i} \geq \frac{p^{2}}{q}
$$

- Step 5: Noting that $r=\left(\sum_{i=1}^{k} q_{i}\right) / p$, we obtain

$$
r \geq \frac{p}{q}
$$

# Mining Data Streams: Introduction 

## Mining Data Streams: Introduction I

- Situation: Data arrives in a stream (or several streams)
- Too much to be put in active storage (main memory, disk, database)
- If not processed immediately or stored (in inaccesible archives), lost forever
- Algorithms involve some summarization of stream(s); e.g.
- create useful samples of stream(s)
- filter the stream(s)
- focus on windows of appropriate length (last $n$ elements)


## Data Streams: Examples

- Sensor data:
- Ocean data (temperature, height): terabytes per day
- Tracking cars (location, speed)
- Image data from satellites
- Internet/web traffic
- Switches that route data also decide on denial of service
- Tracking trends via analyzing clicks


## Data Stream Management System



A data stream management system
Adopted from mmds.org

## Data Stream Queries

- Standing queries
- need to be answered throughout time
- Answers need to be updated when they change
- Example: current or maximum ocean temperature
- Ad-hoc queries
- ask immediate questions
- Example: number of unique users of a web site in the last 4 weeks
- Not all data can be stored/processed Only certain questions feasible
- Need to prepare for queries For example, store data from sliding windows


## Data Stream Queries

## Issues

- Streams deliver elements rapidly: need to act quickly
- Thus, data to work on should fit in main memory
- New techniques required:

Compute approximate, not exact answers
Hashing is a useful technique

## Sampling Elements from a Stream

## SAMPLING ElEMENTS

- Situation:
- Select subsample from stream to store
- Subsample should be representative of stream as a whole
- Running Example:
- Search engine processes stream of search queries
- Stream consists of tuples (user,query,time)
- Can store only $1 / 10$-th of data
- Stream Query: Fraction of repeated search queries?


## Running Example: Pitfall

- Running Example:
- Stream Query: Fraction of repeated search queries?

Naive and bad approach

- For each query, generate random integer from $[0,9]$
- Keep only queries if 0 was generated
- Scenario: Suppose a user has issued
- s queries one time
-d queries two times
- no queries more than two times
- Correct answer is $\frac{d}{d+s}$


## Running Example: Pitfall

- Running Example:
- Stream Query: Fraction of repeated search queries?


## Naive and bad approach

- Correct answer is $\frac{d}{d+s}$
- But on randomly selected queries, we see that
- Of one-time queries, $s / 10$ appear to show once
- Of two-time queries, $d / 10 \times d / 10$ appear to show twice
- Of two-time queries, $d(1 / 10 \times 9 / 10) \times 2$ appear to show once
- Resulting in estimate

$$
\frac{0.01 d}{0.1 s+0.18 d}=\frac{d}{10 s+19 d}
$$

for repeated search queries, which is wrong for positive $s, d$

## Running Example: Pitfall

- Running Example:
- Stream Query: Fraction of repeated search queries?


## Better approach

- For each user (not query!), generate random integer from $[0,9]$
- Keep $1 / 10$ th of users, e.g. if 0 was generated
- Implement randomness by hashing users to 10 buckets
- avoids storing for each user whether he was in or out
- For maintaining sample for $a / b$-th of data, use $b$ buckets, and keep users in buckets 0 to $a-1$


## Running Example: Pitfall

## Better approach

- General Sampling Problem: Generalize from one-valued key to arbitrary-valued keys, keep $a / b$-th of (multi-valued) keys by the same technique
- Reducing sample size: On increasing amounts of data, ratio of data used for sample to be lowered
- When lowering is necessary, decrease $a$ by 1 , so 0 to $a-2$ are still accepted
- Remove all elements with keys hashing to $a-1$


## Filtering Streams

## Filtering Streams: Motivating Example

- Problem: Filter for data for which certain conditions apply
- Can be easy: data are numbers, select numbers of at most 10
- Challenge:
- There is a set $S$ that is too large to fit in main memory
- Condition is too check whether stream elements belong to $S$


## Filtering Streams: Motivating Example

Motivating Example: Email Spam

- Streamed data: pairs (email address, email text)
- Set $S$ is one billion $\left(10^{9}\right)$ approved (no spam!) addresses
- Only process emails from these addresses
need to determine whether arbitrary address belongs to them
- But, addresses cannot be stored in main memory
- Option 1: make use of disk accesses
- Option 2 (preferrable): Devise method without disk accesses, and determine set membership correctly in majority of cases
- Solution: "Bloom Filtering"


## Bloom Filtering: Running Example

- Assume that main memory is 1 GB
- Bloom filtering: use main memory as bit array (of eight billion bits)
- Devise hash function $h$ that hashes email addresses to eight billion buckets
- Hash each member of $S$ (allowed email addresses) to one of the buckets
- Set bits of hashed-to buckets to 1, leave other bits 0
- About $1 / 8$-th of bits are 1


## Bloom Filtering: Running Example

- Hash any new email address:
- If hashed-to bit is 1, classify address as no spam
- If hashed-to bit is 0 , classify address as spam
- Each address hashed to 0 is indeed spam
- But: About $1 / 8$-th of spam emails hash to 1
- So, not each address hashed to 1 is no spam
- $80 \%$ of emails are spam: filtering out $7 / 8$-th is a big deal
- Filter cascade: filter out 7/8-th of (remaining) spam in each step


## Bloom Filter: Definition

## Definition [Bloom Filter]

A Bloom filter consists of

- A bit array $B$ of $n$ bits, initially all zero
- A set $S$ of $m$ key values
- Hash functions $h_{1}, \ldots, h_{k}$ hashing key values to bits of $B$

Number of buckets is $n$


A Bloom filter for set $S=\{x, y, z\}$ using three hash functions

## Bloom Filter: Definition

## Definition [Bloom Filter]

A Bloom filter consists of

- A bit array $B$ of $n$ bits, initially all zero
- A set $S$ of $m$ key values
- Hash functions $h_{1}, \ldots, h_{k}$ hashing key values to bits of $B$

Number of buckets is $n$

## Bloom Filter Workflow

- Initialization
- Take each key value $K \in S$
- Set all bits $h_{1}(K), \ldots, h_{k}(K)$ to one
- Testing keys:
- Take key $K$ to be tested
- Declare $K$ to be a member of $S$ if all $h_{1}(K), \ldots, h_{k}(K)$ are one


## Bloom Filtering: Analysis

- If $K \in S$, all $h_{1}(K), \ldots, h_{k}(K)$ are one, so $K$ passes
- If $K \notin S$, all $h_{1}(K), \ldots, h_{k}(K)$ could be one, so $K$ mistakenly passes False positive!
- Goal: Calculate probability of false positives
- For that, calculate probability that bit is zero after initialization
- Relates to throwing $y$ darts at $x$ targets, where
- Targets are bits in array, so $x=n$
- Darts are members in $S(=m)$ times hash functions $(=k)$, which makes $y=k m$

What is the probability that target is not hit by any dart?

## Bloom Filtering: Analysis

Throwing $y$ darts at $x$ targets:

- Probability that a given dart will not hit a given target is $(x-1) / x$
- Probability that none of the $y$ darts will hit a given target is

$$
\begin{equation*}
\left(\frac{x-1}{x}\right)^{y}=\left(1-\frac{1}{x}\right)^{x \frac{y}{x}} \tag{2}
\end{equation*}
$$

- By $(1-\epsilon)^{1 / \epsilon}=1 / e$ for small $\epsilon$, we obtain that (2) is $e^{-y / x}$
- $x=n, y=k m$ : probability that a bit remains 0 is $e^{-k m / n}$
- Would like to have fraction of 0 bits fairly large
- If $k$ is about $n / m$, then probability of a 0 is $e^{-1}$ (about $37 \%$ )
- In general, probability of false positive is $k 1$ bits, which evaluates as

$$
\begin{equation*}
\left(1-e^{-\frac{k m}{n}}\right)^{k} \tag{3}
\end{equation*}
$$

## Counting Distinct Elements The Flajolet-Martin Algorithm

## Counting Distinct Elements: Problem

- Problem: Elements in streams can be identical
- Question: How many different elements has the stream brought along?
- Model: Consider the universal set of all possible elements
- Consider stream as a subset of the universal set
- Question becomes: What is the cardinality (size) of this subset?
- Example: Unique users of website
- Amazon: determine number of users from user logins
- Google: determine number of users from search queries


## Counting Distinct Elements: Problem

- Situation: Stream picks elements from universal set
- Question: Size of subset of elements appearing in stream?
- Obvious, but expensive:
- Keep stream elements in main memory
- Store them in efficient search structure (hash table, search tree)
- Works for sufficiently small amounts of distinct elements
- If too many distinct elements, or too many streams:
- Use more machines Ok if affordable
- Use secondary memory (disk) slow
- Here: Estimate number of distinct elements using much less main memory than needed for storing all distinct elements
- The Flajolet-Martin algorithm does this job


## The Flajolet-Martin Algorithm

- Central idea: Hash elements to bit strings of sufficient length
- For example, to hash URL's, 64-bit strings are sufficiently long
- Intuition:
- The more different elements, the more different bit strings
- The more different bit strings, the more "unusual" bit strings
- Unusual here = bit string starts with many zeroes


## Definition [TAIL LengTh]

- Let $h$ be the hash function that hashes stream elements $a$ to bit strings $h(a)$
- The tail length of $h(a)$ is the number of zeroes by which it begins


## The Flajolet-Martin Algorithm

## Definition [TAIL Length]

- Let $h$ be the hash function that hashes stream elements $a$ to bit strings $h(a)$
- The tail length of $h(a)$ is the number of zeroes by which it begins
- Alternatively: $h(a)$ number of zeroes a string ends with


## Flajolet Algorithm

- Let $A$ be the set of stream elements
- Let

$$
\begin{equation*}
R:=\max _{a \in A} h(a) \tag{4}
\end{equation*}
$$

be the maximum tail length observed among stream elements

- Estimate $2^{R}$ for the number of distinct elements in the stream


## Flajolet-Martin Algorithm: Example



Hashing user names to 8-bit strings
From towardsdatascience.com

## Flajolet-Martin Algorithm: Explanation

- Probability that bit string $h(a)$ starts with $r$ zeroes is $2^{-r}$
- Probability that none of $m$ distinct elements has tail length at least $r$ is

$$
\begin{equation*}
\left(1-2^{-r}\right)^{m}=\left(\left(1-2^{-r}\right)^{2^{r}}\right)^{m 2^{-r}(1-\epsilon)^{1 / \epsilon} \approx 1 / e}=e^{-m 2^{-r}} \tag{5}
\end{equation*}
$$

- Let $P_{m, r}:=1-\left(1-2^{-r}\right)^{m} \approx 1-e^{-m 2^{-r}}$ be the probability that for $m$ stream elements, the maximum tail length $R$ observed is at least $r$.
- Conclude:
- For $m \gg 2^{r}$, it holds that $P_{m, r}$ approaches 1
- For $m \ll 2^{r}$, it holds that $P_{m, r}$ approaches 0
- So, $2^{R}$ is unlikely to be much larger or much smaller than $m$


## Flajolet-Martin Algorithm: Combining

## EsTIMATES

- Idea: Use several hash functions $h_{k}, k=1, \ldots, K$
- Combine their estimates $X_{k}, k=1, \ldots, K$
- Pitfall 1: Averaging
- Let $p_{r}$ be the probability that the maximum tail length of $h_{k}$ is $r$
- One can compute that

$$
p_{r} \geq 2 p_{r-1} \geq \ldots \geq 2^{-r+1} p_{1} \geq 2^{-r} p_{0}
$$

- So $E\left(X_{k}\right)$, the expected value of $X_{k}$ computes as

$$
E\left(X_{k}\right)=\sum_{r \geq 0} p_{r} 2^{r} \geq p_{0} \sum_{r \geq 0} 2^{-r} 2^{r}=p_{0} \sum_{r \geq 0} 1=\infty
$$

- Therefore $\frac{1}{K} \sum_{k=1}^{K} E\left(X_{k}\right)$ the expected value of the average of the $X_{k}$ turns out to be infinite as well
- Conclusion: Overestimates spoil averaging


## Flajolet-Martin Algorithm: Combining

## ESTIMATES

- Idea: Use several hash functions $h_{k}, k=1, \ldots, K$
- Combine their estimates $X_{k}, k=1, \ldots, K$
- Pitfall 2: Computing Medians
- The median is always a power of two makes only very limited sense
- Solution:
- Group the hash functions into small groups and take averages within groups
- Estimate $m$ as median of group averages
- Groups should be of size $C \log _{2} m$ for some small $C$
- Space Requirements: Need to store only value of $X_{k}$, requiring little space as a maximum


## Materials / Outlook

- See Mining of Massive Datasets: section 2.6; sections 4.1-4.4
- As usual, see http://www.mmds.org/in general for further resources
- For deepening your understanding, consider voluntary homework: read 2.6.7 and try to make sense of this
- Next lecture: "Mining Data Streams II / PageRank I"
- See Mining of Massive Datasets 4.5-4.7; 5.1-5.2

