



# Time Series Databases and Streaming algorithms



# Introduction and motivation for Time Series



# Internet of things



# Domotics





# Predictive Maintenance



# Environmental tracking

*CarbonTracker* free troposphere CO<sub>2</sub>  
2008-Jun-06





A time series is a sequence of data points,  
typically consisting of successive measurements  
made over a time interval.





# Why Time Series Databases?

- High Volume of Data
- Large Quantities of Immutable Data
- Is Primarily Sorted Temporally
- Needs to Be Rolled Up To Gain Majority of Insights
- Needs to Be Normalized Across Multiple Time Zones



# Problems using Relational DBs

1. It's Difficult to Change the Sample Rate
2. It's Difficult To Use SQL Queries For Analysis
3. Time Zones Add Extra Complexity To Your Data Analysis



# Advantages of NoSQL

1. Greater simplicity in the DB engine
2. Ability to handle semi-structured and denormalized data
3. Potentially much higher scalability



# Disadvantages of NoSQL

1. Higher complexity in the application
2. Loss of abstraction provided by the query optimizer



# Basic Operations on Time Series Data



# What do we need to do with TS

- Acquire
  - Measurement, transmission, reception
- Store
- Retrieve
- Analyze and visualize



# Rescaling

- Transform the range of variation to a given scale
- Useful for algorithms sensitive to the magnitude of the signal



# Resampling

- Differences in sampling resolution
- Bring both series to the same sample frequency
- Requires a function for collapsing points together





# Shifting

- Align series we know are misaligned
- Bad reference time, drifting clock, ...



# Slicing

- Retrieve a time series based on a given time range



# Dynamic Time Warping

- Used for measuring similarity between series that vary in time or speed
- Dynamic time warping is a sequence alignment technique used in speech recognition
- It is an algorithm that has  $O(n^2)$  complexity



# Subsequence Matching

- A sequence query is matched against a longer TS
- Also related with Chunking where we look for repeating patterns



# Statistical measures

- Mean
- Median
- Standard Deviation
- Variance
- Quantiles



# Statistical fitting

- Interpolation
- Linear models
- Non linear models



# Data Storage for Time Series Data

# Log Files

- Simplest solution
- Right solution when low number of time series or data fits in memory

1950	1	0.920000E+00
1950	2	0.400000E+00
1950	3	-0.360000E+00
1950	4	0.730000E+00
1950	5	-0.590000E+00
1950	6	-0.600000E-01
1950	7	-0.126000E+01
1950	8	-0.500000E-01
1950	9	0.250000E+00
1950	10	0.850000E+00
1950	11	-0.126000E+01
1950	12	-0.102000E+01
1951	1	0.800000E-01
1951	2	0.700000E+00
1951	3	-0.102000E+01
1951	4	-0.220000E+00
1951	5	-0.590000E+00
1951	6	-0.164000E+01
1951	7	0.137000E+01
1951	8	-0.220000E+00
1951	9	-0.136000E+01
1951	10	0.187000E+01



# Advanced Log Files

- Same concept about storing TS in files
- Use a smart binary encoding format
- Allows less processing, aka no parsing
- Stores data more efficiently for scan readings

```
message AddressBook {  
    required string owner;  
    repeated string ownerPhoneNumbers;  
    repeated group contacts {  
        required string name;  
        optional string phoneNumber;  
    }  
}
```



# Advanced Log Files

- Lots of binary formats lately
  - Thrift
  - Avro
  - Parquet

*We created Parquet to make the advantages of compressed, efficient columnar data representation available to any project in the Hadoop ecosystem.*

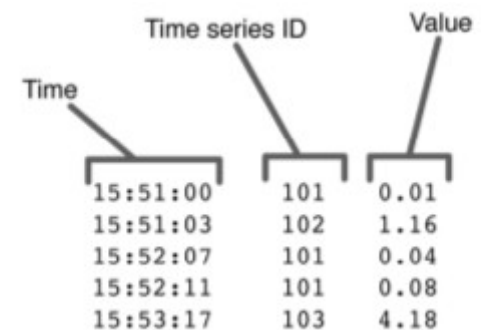


# Relational Databases

- True and tested technology validated in multitude of scenarios
- Allows indexing out of the box
- Allows data replication and sharding (to some extent)

# Relational Databases

- Use the Star Schema
- The fact table contains the measurements
- The dimension tables contains info about the series



The diagram shows a table with three columns: Time, Time series ID, and Value. Each column has a bracket above it with a label pointing to it: 'Time' for the first column, 'Time series ID' for the second, and 'Value' for the third. The table contains five rows of data.

Time	Time series ID	Value
15:51:00	101	0.01
15:51:03	102	1.16
15:52:07	101	0.04
15:52:11	101	0.08
15:53:17	103	4.18



# Relational Databases

- The Star Schema can work reasonably to the hundreds of millions
- We can even implement the Star Schema in a NoSQL database
- When data grows this size several problems arise mostly related to the Star Schema itself.



# Limitations of the Star Schema

- It uses one row per measurement
- Limitants of retrieval speed:
  - number of rows scanned,
  - total number of values retrieved
  - total volume of data retrieved



# NoSQL databases

- Most of TS DBs use a NoSQL engine
  - OpenTSB → Hbase
  - InfluxDB → BoltDB
  - Prometheus → LevelDB
  - Newts → Cassandra

# NoSQL databases

- Tall and narrow vs Short and wide table designs
- Short and wide denormalizes data
- Short and wide provides several advantages over the columnar data model

Time series ID	Time-window start time	Columns are named by sample time offset within time window				
		+0	+3	+7	+11	+17
101	15:51:00	0.01				
101	15:52:00			0.04	0.08	
102	15:51:00		1.16			
103	15:53:00					4.18



# NoSQL databases


- Indexed by TS and timestamp the most common access pattern
- Retrieving data is an almost sequential reading from disk

The diagram illustrates a time series data table structure. It includes a table with four rows and columns for Time series ID, Time-window start time, and data values at different offsets. Annotations explain the components: 'Time series ID' points to the first column, 'Time-window start time' points to the second column, 'Columns are named by sample time offset within time window' points to the header row of the data columns, 'Row key' points to the first two columns, and 'Data values' points to the data columns.

Time series ID	Time-window start time	Columns are named by sample time offset within time window				
		+0	+3	+7	+11	+17
101	15:51:00	0.01				
101	15:52:00			0.04	0.08	
102	15:51:00		1.16			
103	15:53:00					4.18


# Improvements over the Wide Table Design

- Collapse all the data into a blob
- Compress the blob so less data has to be read
- Allow coexistence of wide table columns and the blob



# Improvements over the Wide Table Design

- Avoid the reads in order to overcome insert bottlenecks
- Create a fallback system in order to prevent failures
- Allow access to the in-memory data






# Why not with RDBMs?

- Why use a RDBMs when you're not using any of its strong points?
- Also some features, ie. transactions, get in your way for scaling

# Time Series Databases





InfluxDB



# Features

- Written in Go
- Using BoltDB as its internal storage engine
- SQL-like language
- HTTP(S) API for querying data
- Stores metrics and event data
- Horizontally scalable



# Key concepts

A **series** is a collection of data **points** along a **timeline** that share a common **key**, expressed as a **measurement** and **tag set** pairing, grouped under a **retention policy**

[https://influxdb.com/docs/v0.9/concepts/key\\_concepts.html](https://influxdb.com/docs/v0.9/concepts/key_concepts.html)





# Key concepts

- **Measurement**
  - It is the value being recorded
  - Can be shared amongst many series
  - All series under a given measurement have the same field keys and differ only in their tag set

[https://influxdb.com/docs/v0.9/concepts/key\\_concepts.html](https://influxdb.com/docs/v0.9/concepts/key_concepts.html)



# Key concepts

- **Tag**
  - It is a key-value pair.
  - A measurement could have several tags
  - Tags are indexed
  - Both the key and value are strings

[https://influxdb.com/docs/v0.9/concepts/key\\_concepts.html](https://influxdb.com/docs/v0.9/concepts/key_concepts.html)



# Key concepts

- **Point**

- A point is a single collection of fields in a series.
- It is uniquely identified by its series and timestamp

[https://influxdb.com/docs/v0.9/concepts/key\\_concepts.html](https://influxdb.com/docs/v0.9/concepts/key_concepts.html)



# Key concepts

- **Field**
  - A field is a key-value pair
  - It records an actual metric for a given point
  - They are not indexed
  - They are required at least 1 on each point

[https://influxdb.com/docs/v0.9/concepts/key\\_concepts.html](https://influxdb.com/docs/v0.9/concepts/key_concepts.html)



# Key concepts

- **Database**
  - similar in concept to RDBS groups series
- **Retention policy**
  - defines what to do with data that is older than the prescribed retention policy

[https://influxdb.com/docs/v0.9/concepts/key\\_concepts.html](https://influxdb.com/docs/v0.9/concepts/key_concepts.html)

# Logging points into InfluxDB

```
{
  "database": "mydb",
  "points": [
    {
      "measurement": "cpu_load",
      "tags": {
        "host": "server01",
        "core": "0"
      },
      "time": "2009-11-10T23:00:00Z",
      "fields": {
        "value": 0.45
      }
    },
    {
      "measurement": "cpu_load",
      "tags": {
        "host": "server01",
        "core": "1"
      },
      "time": "2009-11-10T23:00:00Z",
      "fields": {
        "value": 1.56
      }
    }
  ]
}
```



# HTTP endpoint

/query GET

/write OPTIONS

/write POST

/ping GET

/ping HEAD

/data/process\_continuous\_queries POST

# Query exploration

Queries like in RDBMs

```
SELECT * FROM cpu
WHERE cpu_1 = '1'
```

Querying by time

```
SELECT mean(value) FROM cpu
WHERE time > 12345678s
GROUP BY time(10m);
SELECT mean from "hour_summaries".cpu
WHERE time > now() - 7d
```



# Dealing with Time

- Querying using time strings

```
SELECT value FROM response_times  
WHERE time > '2013-08-12 23:32:01.232' and time < '2013-08-13';
```

- Relative time

```
SELECT value FROM response_times  
WHERE time > now() - 1h limit 1000;
```

- Absolute time

```
SELECT value FROM response_times  
WHERE time > now() - 1h limit 1000;
```

# Dealing with missing values

- Use null, previous, none for missing values

```
SELECT COUNT(type) FROM events  
WHERE time > now() - 3h  
GROUP BY time(1h) fill(null)
```

```
SELECT COUNT(type) FROM events  
WHERE time > now() - 3h  
GROUP BY time(1h) fill(previous)
```

```
SELECT COUNT(type) FROM events  
WHERE time > now() - 3h  
GROUP BY time(1h) fill(none)
```



# Write data

- Ingest data into InfluxDB using the HTTP API

- Create the Database

```
curl -G http://localhost:8086/query --data-urlencode "q=CREATE  
DATABASE mydb"
```

- Write data into the database

```
curl -i -XPOST 'http://localhost:8086/write?db=mydb' --data-binary  
'cpu_load_short,host=server01,region=us-west value=0.64  
1434055562000000000'
```



# Hands on

- Import data from Standard&Poor
- Explore the performance of different encodings:
  - Several fields for a single point
  - Each column as a separate TS
- Create the following queries:
  - Select maximum opening price on a given period for each quote
  - Select the monthly average



# Hands on (Advanced)

- Import extra dataset
- Compare loading and querying data between MySQL and InfluxDB



# Streaming data




Algorithms for processing data streams in which  
the input is presented as a sequence of items  
and can be examined in only a few passes



# Examples



# Examples: Anomaly Detection



# Real Time Telemetry








# Streaming algorithms




# Characteristics of streaming algorithms

- Operates on a continuous stream of data
- Unknown or infinite size
- Only one pass, that allows following options:
  - Store it
  - Lose it
  - Store an approximation of it
- Limited processing time per item
- Limited total memory



These algorithms produce an approximate answer based on a summary or "sketch" of the data stream in memory



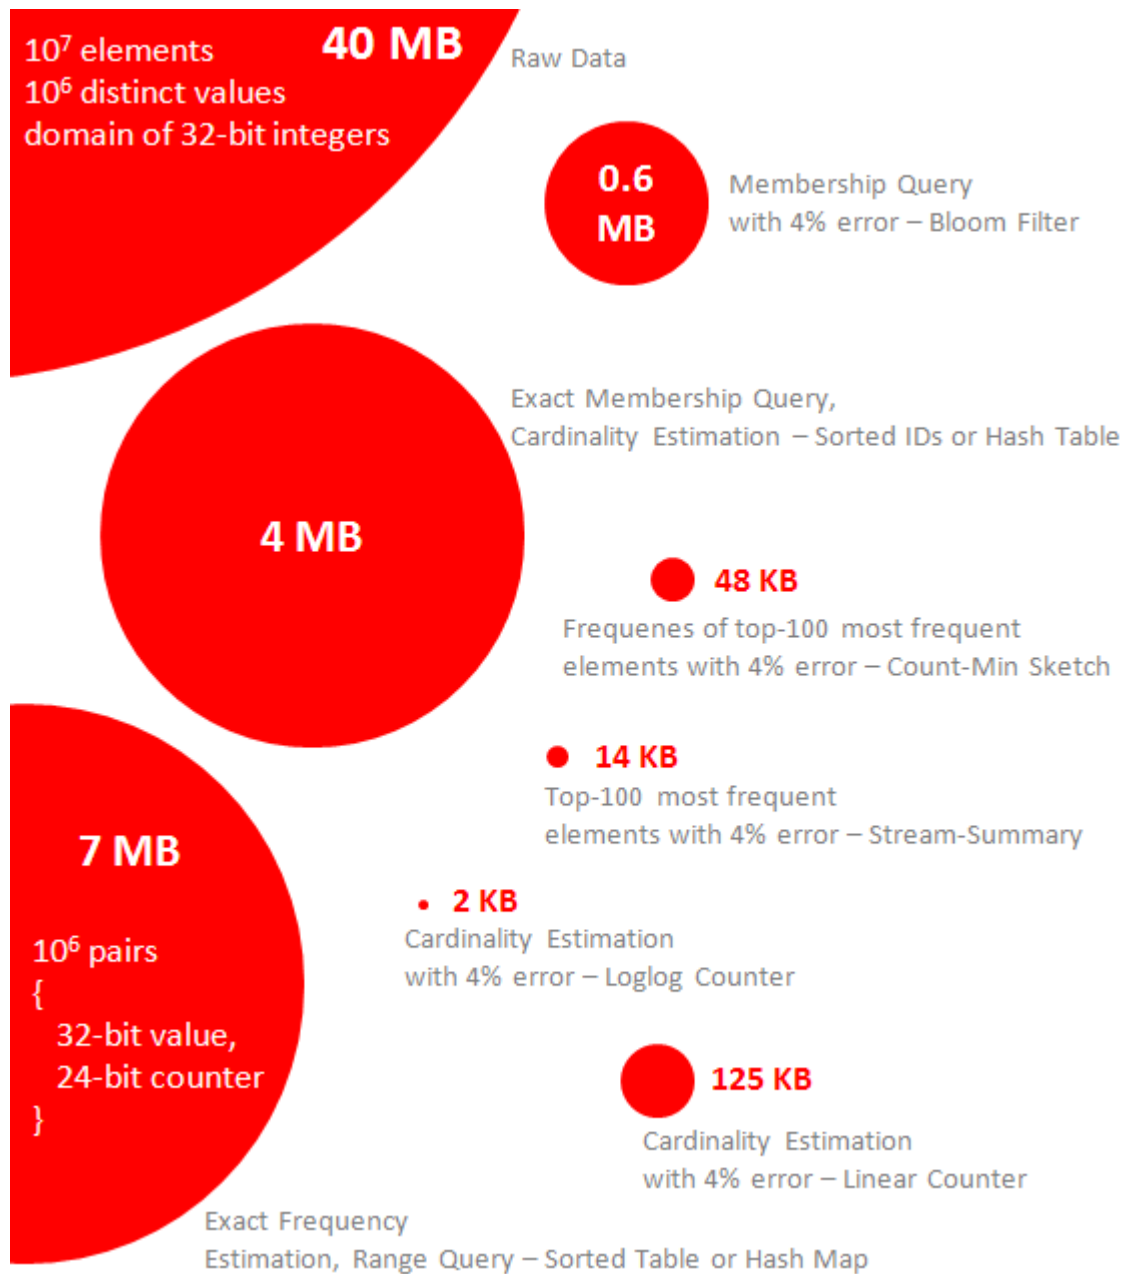
They have limited memory available to them  
(much less than the input size) and also  
limited processing time per item.



# Questions to answer

- Frequency moments
- Counting distinct elements
- Heavy Hitters
- Anomaly detection / Membership query
- Online learning





# Cardinality estimation

## Linear Counting

```
1  class LinearCounter {
2      BitSet mask = new BitSet(m) // m is a design parameter
3
4      void add(value) {
5          int position = hash(value) // map the value to the range 0..m
6          mask.set(position) // sets a bit in the mask to 1
7      }
8  }
```

Load Factor is the ratio of distinct elements over the size  $m$

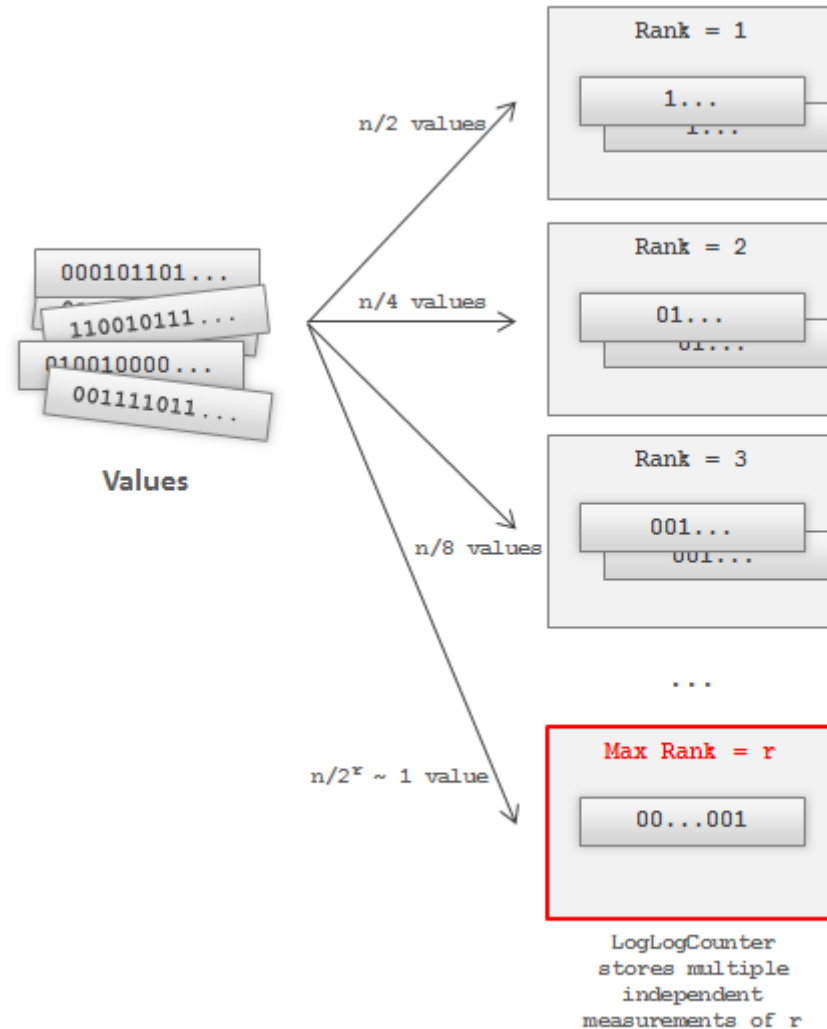
# Cardinality estimation

## Linear Counting

$\hat{n} = -m \ln \frac{m - w}{m}$	<p><math>\hat{n}</math> - cardinality estimation <math>w</math> - mask weight (a number of 1's) <math>m</math> - mask size</p>
$bias(\frac{\hat{n}}{n}) = E(\frac{\hat{n}}{n}) - 1 = \frac{e^t - t - 1}{2n}$	<p>This equation expresses a bias of the estimation (the ratio between estimation and true cardinality) as a function of the load factor and expected cardinality (or upper bound). <math>t</math> - load factor, <math>n/m</math> <math>E(.)</math> - mathematical expectation <math>n</math> - maximum cardinality (or upper bound, or capacity)</p>
$m > \max(5, 1/(\varepsilon t)^2) \cdot (e^t - t - 1)$	<p>A practical formula that allow one to choose <math>m</math> by the standard error of the estimation. <math>m</math> - mask size <math>\varepsilon</math> - standard error of the estimation <math>t</math> - load factor, <math>n/m</math></p>

# Cardinality estimation

## Loglog Counting




# Cardinality estimation

## Loglog Counting

```
1  class LogLogCounter {
2      int H          // H is a design parameter
3      int m = 2^k    // k is a design parameter
4      etype[] estimators = new etype[m] // etype is a design parameter
5
6      void add(value) {
7          hashedValue = hash(value)
8          bucket = getBits(hashedValue, 0, k)
9          estimators[bucket] = max(
10             estimators[bucket],
11             rank( getBits(hashedValue, k, H) )
12         )
13     }
14
15     getBits(value, int start, int end)
16     rank(value)
17 }
```

# Frequency Estimation: Count-Min Sketch



# Frequency Estimation: Count-Min Sketch

```
1  class CountMinSketch {
2      long estimators[][] = new long[d][w]    // d and w are design parameters
3      long a[] = new long[d]
4      long b[] = new long[d]
5      long p    // hashing parameter, a prime number. For example 2^31-1
6
7      void initializeHashes() {
8          for(i = 0; i < d; i++) {
9              a[i] = random(p)    // random in range 1..p
10             b[i] = random(p)
11         }
12     }
13
14     void add(value) {
15         for(i = 0; i < d; i++)
16             estimators[i][ hash(value, i) ]++
17     }
18
19     long estimateFrequency(value) {
20         long minimum = MAX_VALUE
21         for(i = 0; i < d; i++)
22             minimum = min(
23                 minimum,
24                 estimators[i][ hash(value, i) ]
25             )
26         return minimum
27     }
28
29     hash(value, i) {
30         return ((a[i] * value + b[i]) mod p) mod w
31     }
32 }
```

# Frequency Estimation: Count-Mean-Min Sketch

```
1  class CountMeanMinSketch {
2      // initialization and addition procedures as in CountMinSketch
3      // n is total number of added elements
4
5      long estimateFrequency(value) {
6          long e[] = new long[d]
7          for(i = 0; i < d; i++) {
8              sketchCounter = estimators[i][ hash(value, i) ]
9              noiseEstimation = (n - sketchCounter) / (w - 1)
10             e[i] = sketchCounter - noiseEstimator
11         }
12         return median(e)
13     }
14 }
```



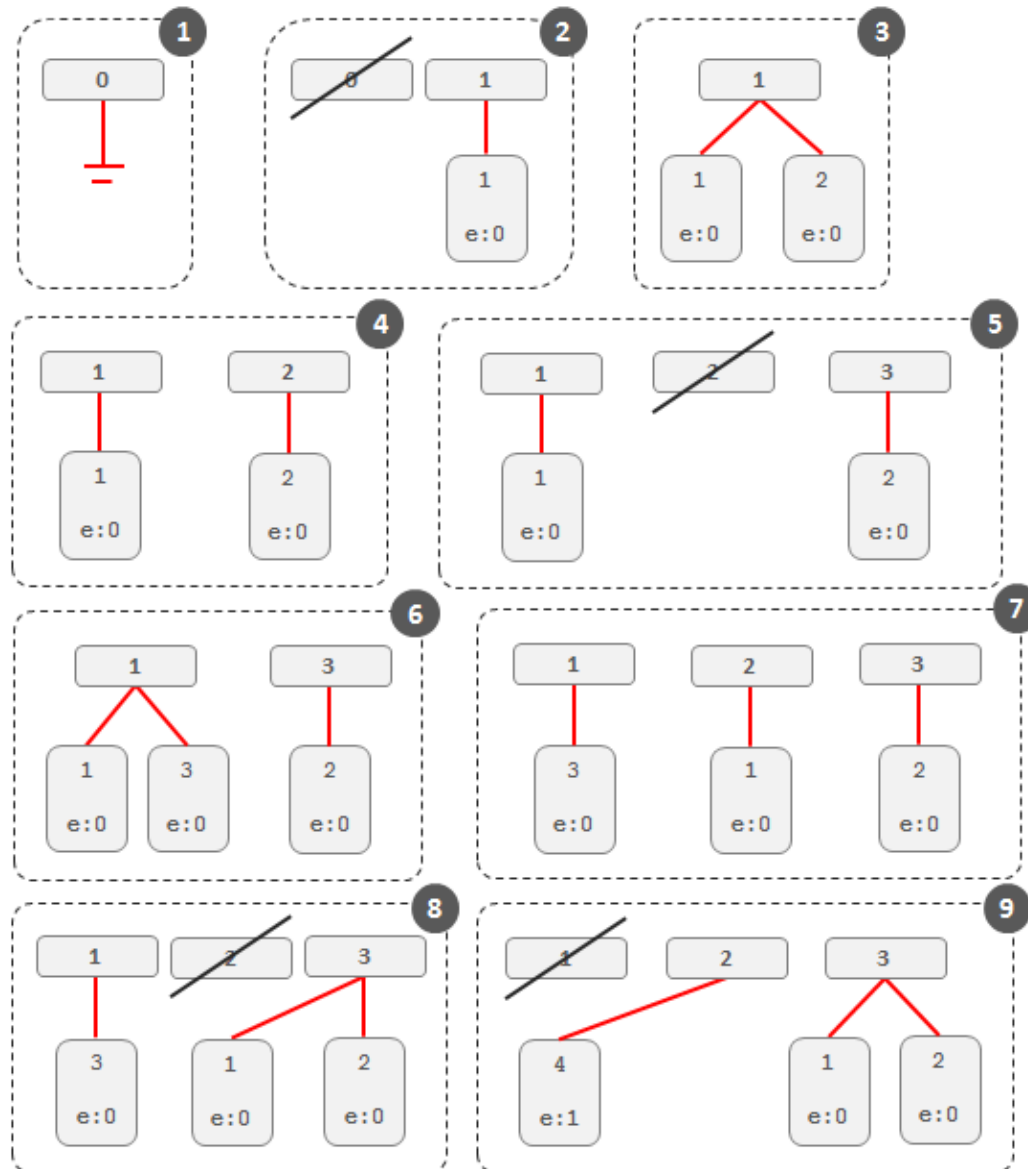
# Heavy Hitters

## Count-Min Sketch


```
1  class CountMeanMinSketch {
2      // initialization and addition procedures as in CountMinSketch
3      // n is total number of added elements
4
5      long estimateFrequency(value) {
6          long e[] = new long[d]
7          for(i = 0; i < d; i++) {
8              sketchCounter = estimators[i][ hash(value, i) ]
9              noiseEstimation = (n - sketchCounter) / (w - 1)
10             e[i] = sketchCounter - noiseEstimator
11         }
12         return median(e)
13     }
14 }
```

# Heavy Hitters

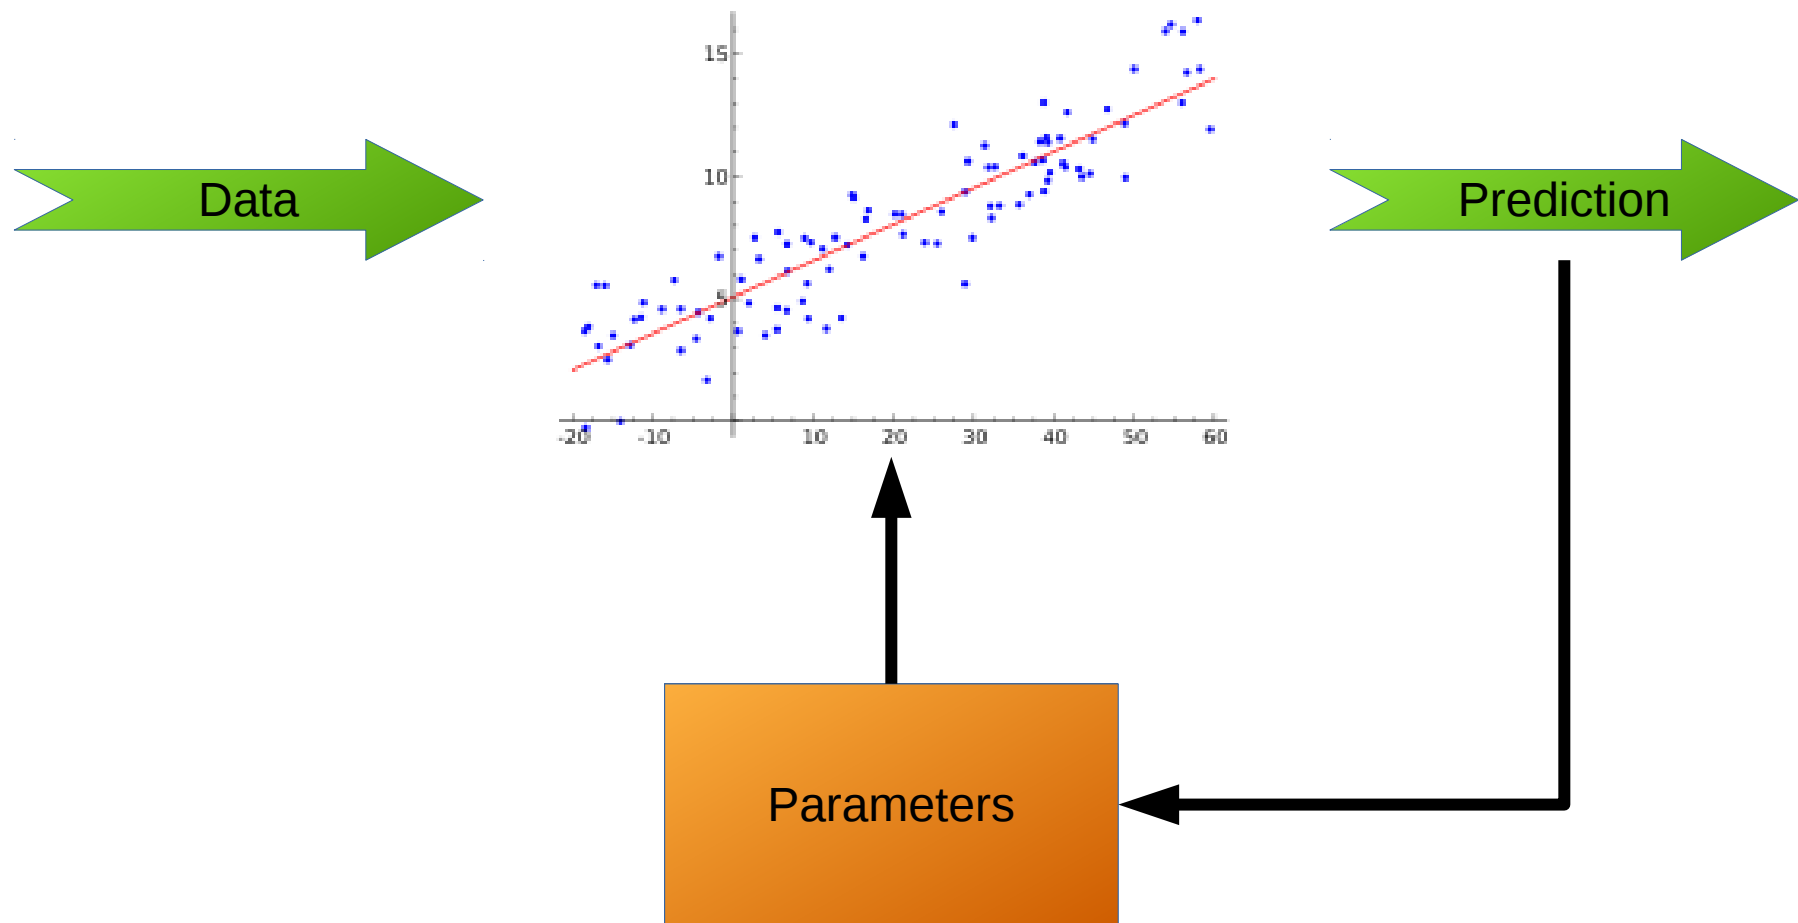
## Stream-Summary



# Membership Query Bloom Filter



# Online Learning



# Feature Hashing

- John likes to watch movies.
- Mary likes movies too.
- John also likes football.

$$\begin{pmatrix} \text{John} & \text{likes} & \text{to} & \text{watch} & \text{movies} & \text{Mary} & \text{too} & \text{also} & \text{football} \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}$$

# Feature Hashing

```
function hashing_vectorizer(features : array of string, N : integer):  
    x := new vector[N]  
    for f in features:  
        h := hash(f)  
        x[h mod N] += 1  
    return x
```

Can be extended to use signed hashing functions


Pros:

- Extremely fast
- No memory footprint

Cons

- There is no way to reverse features


# Stochastic Gradient Descent





# Apache Spark





**Program  
Model**



**Storage  
System**

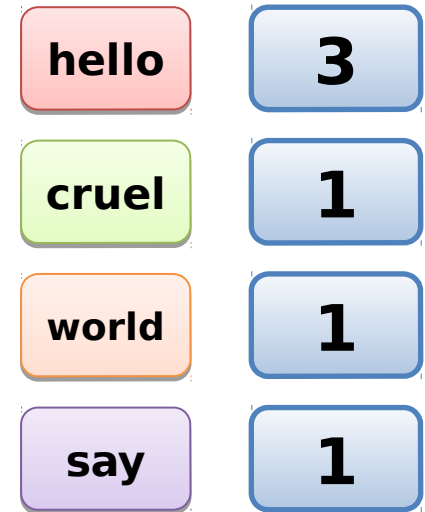
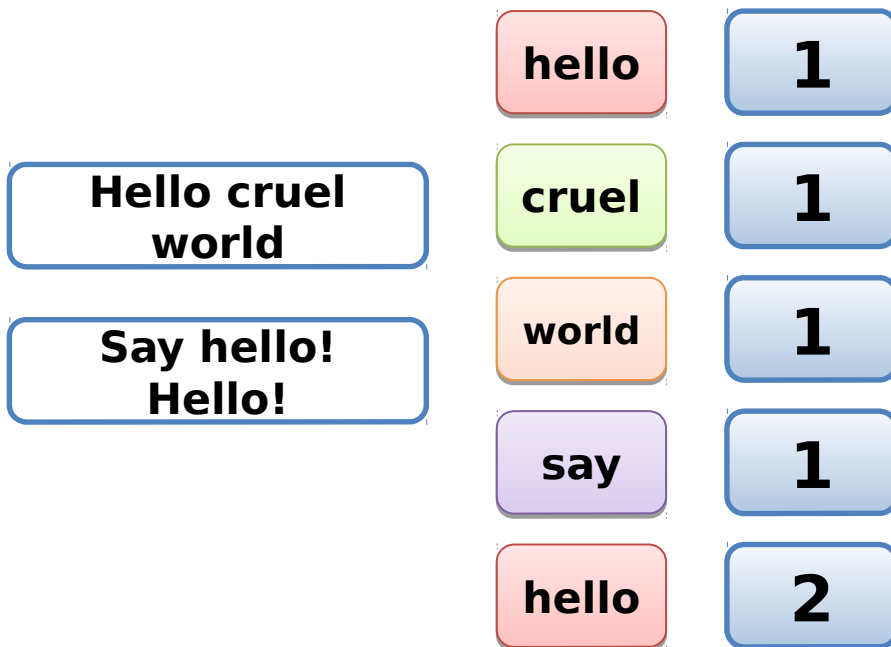
# Word Count

Raw


Map

Reduce

Result



# Problem with Iterative Algos



Disk I/O is very expensive

# Opportunity for a new approach

- Keep data in memory
- Use a new distribution model





# Spark Streaming



# Resilient Distributed Dataset (RDDs)


- A distributed and immutable collection of objects
- Each RDD can be split into multiple partitions
- RDDs allow two types of operations:
  - Transformations (lazy)
  - Actions (non-lazy)



# DStream


A sequence of RDDs representing a stream of data

# DStreams

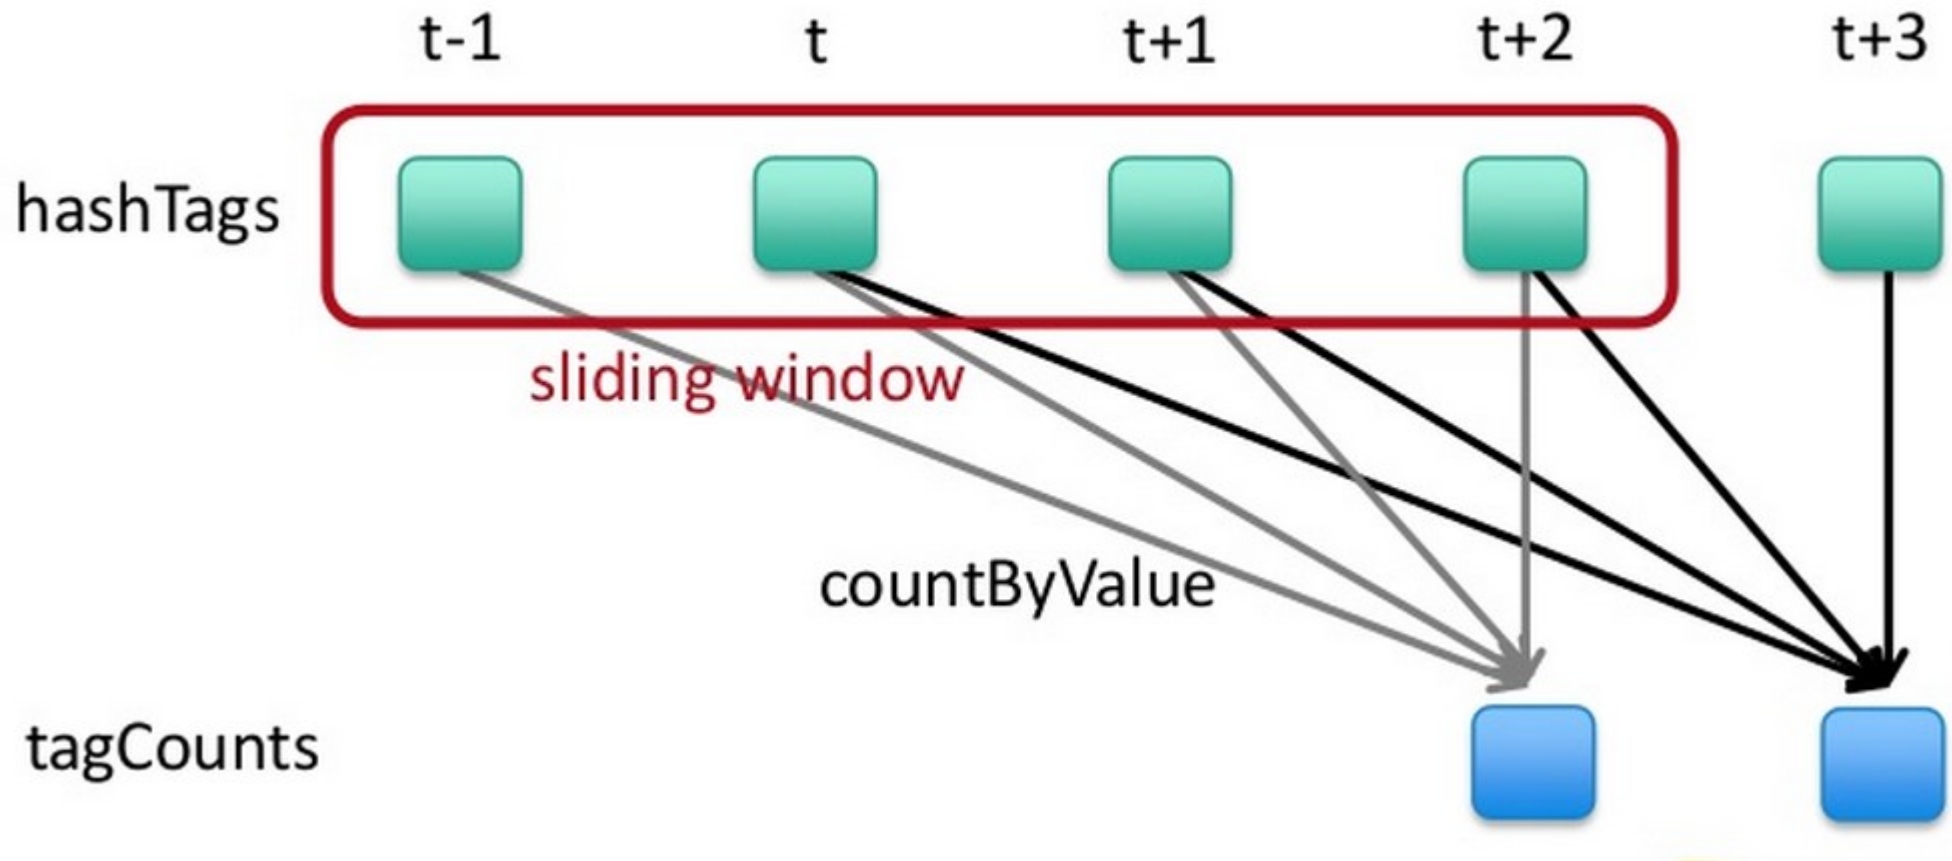





# Windows



# Windowing computations



# Stateful computations





# DStream API

<http://spark.apache.org/docs/1.3.1/api/scala/index.html#org.apache.spark.streaming.dstream.DStream>



# Hands on Streaming

- Start the Spark server
- Create a Job
- Run netcat
- Send



# Hands on Streaming (Advanced)

- Implement Count-Log on basic Spark

# Setup environment

- Prerequisite:
  - Install latest version of Vagrant
  - Install latests version of Virtualbox

<https://www.vagrantup.com/>

<https://www.virtualbox.org/>

- Create the Virtual Machine:

```
vagrant init codezomb/trusty64-docker
```

```
vagrant up
```

<http://blog.scottlowe.org/2015/02/10/using-docker-with-vagrant/>

# Setup environment

- Log in into the VM machine

```
vagrant ssh
```

- Install some Ubuntu packages

```
sudo apt-get update
```

```
sudo apt-get -y install docker
```

```
openjdk-7-jdk
```

- Pull docker images

```
docker pull tutum/influxdb
```

<http://old.blog.phusion.nl/2013/11/08/docker-friendly-vagrant-boxes/>

```
docker pull sequenceiq/spark:1.3.0
```





@tonicebrian, @Enerbyte



[toni.cebrian@gmail.com](mailto:toni.cebrian@gmail.com)



<https://es.linkedin.com/in/tonicebrian>



<http://www.tonicebrian.com>

**WE'RE  
HIRING!**