How *t*-digest works and why

Ted Dunning

June 1, 2015

T-digest

Ted Dunning, Chief Applications Architect MapR Technologies

Email tdunning@mapr.com tdunning@apache.org

Twitter @Ted_Dunning

A New Look at Anomaly Detection

by Ted Dunning and Ellen Friedman © June 2014 (published by O'Reilly)

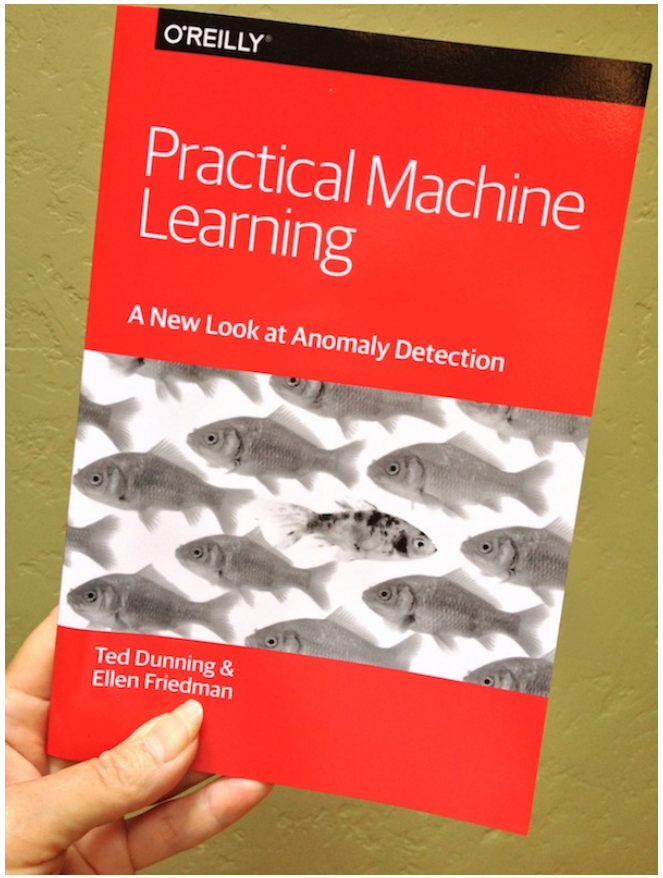

e-book available courtesy of MapR

http://bit.ly/1jQ9QuL

Last October: *Time Series Databases*

by Ted Dunning and Ellen Friedman © Oct 2014 (published by O'Reilly)

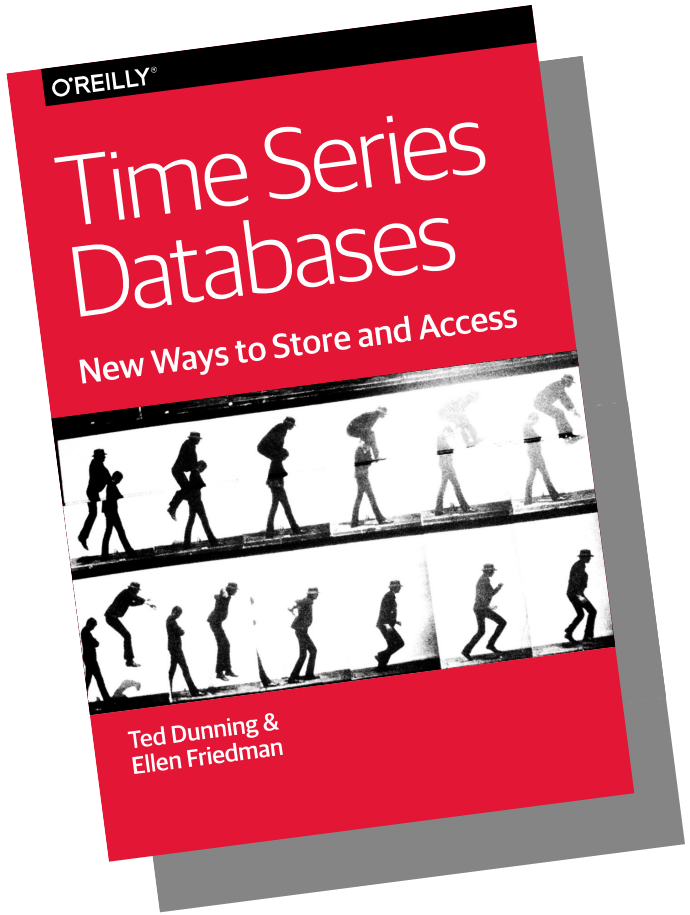

Available Now: *Real World Hadoop*

by Ted Dunning and Ellen Friedman © Feb 2015 (published by O'Reilly)

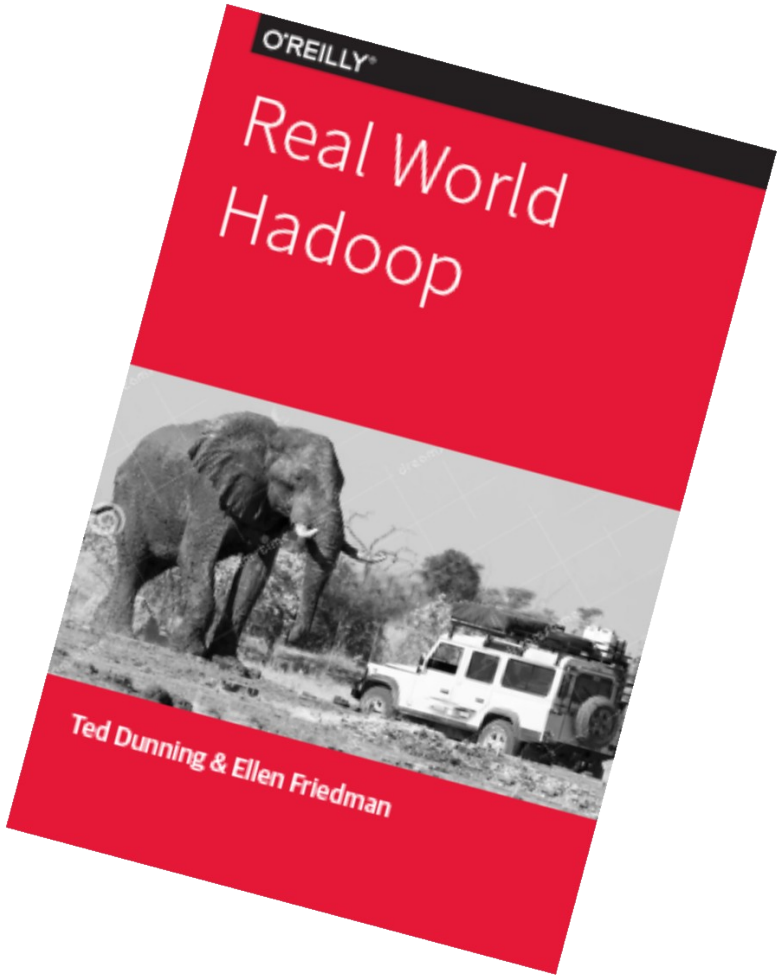

Practical Machine Learning series (O'Reilly)

- Machine learning is becoming mainstream
- Need pragmatic approaches that take into account real world business settings:
  - Time to value
  - Limited resources
  - Availability of data
  - Expertise and cost of team to develop and to maintain system
- Look for approaches with big benefits for the effort expended

Agenda

- Why should we estimate quantiles?
- How t-digest works
- How can you get it?
- Questions

Why on-line algorithms?

7 Clicks on 10 Links - Jun 1, 2015

16	**berlinbuzzwords.de/15/schedule** berlinbuzzwords.de · Created May 20 11:47 PM ·
9	**meetup.com/SF-Bay-Areas-Big-Data-Think-Tank/events/222536108** meetup.com · Created May 20 12:04 AM ·
5	**bigdataeverywhere.com/boston-hadoop-conference** bigdataeverywhere.com · Created May 30 5:10 PM ·
4	**mapr.com/practical-machine-learning** mapr.com · Created Feb 15, 2014 1:35 AM ·

Why Quantiles (percentiles)

Suppose You Have …

- 100 M users, 1K sites touched each day
- What is 99.9% latency for each user/site combination?
  - for each user?
  - for each site?
  - for users in Kansas?
  - for users who complained?
  - for users who complained, but before they complained?

Or Suppose …

- 1000 nodes, each with 24 disks, 100 unique RPC calls
- Want latencies for all disks, all RPC calls between all nodes
  - 50 %-ile, 99%-ile, 99.9%-ile
- <100ns overhead per measurement
- <10MB overhead per node
- No logs except for exceptionally slow cases
- Summary at any time

What about accuracy?

What Accuracy Required?

- 50%-ile ± 0.5%
- 99.99%-ile ± 0.5%
- 99.99%-ile ± 0.001%
- 50%-ile ± 0.001%

What Accuracy Required?

- 50%-ile ± 0.5% — Often just fine
- 99.99%-ile ± 0.5% — *Nonsense*
- 99.99%-ile ± 0.001% — By definition
- 50%-ile ± 0.001% — *Over-kill*

The internals

Variable Cluster Size for Constant Relative Accuracy

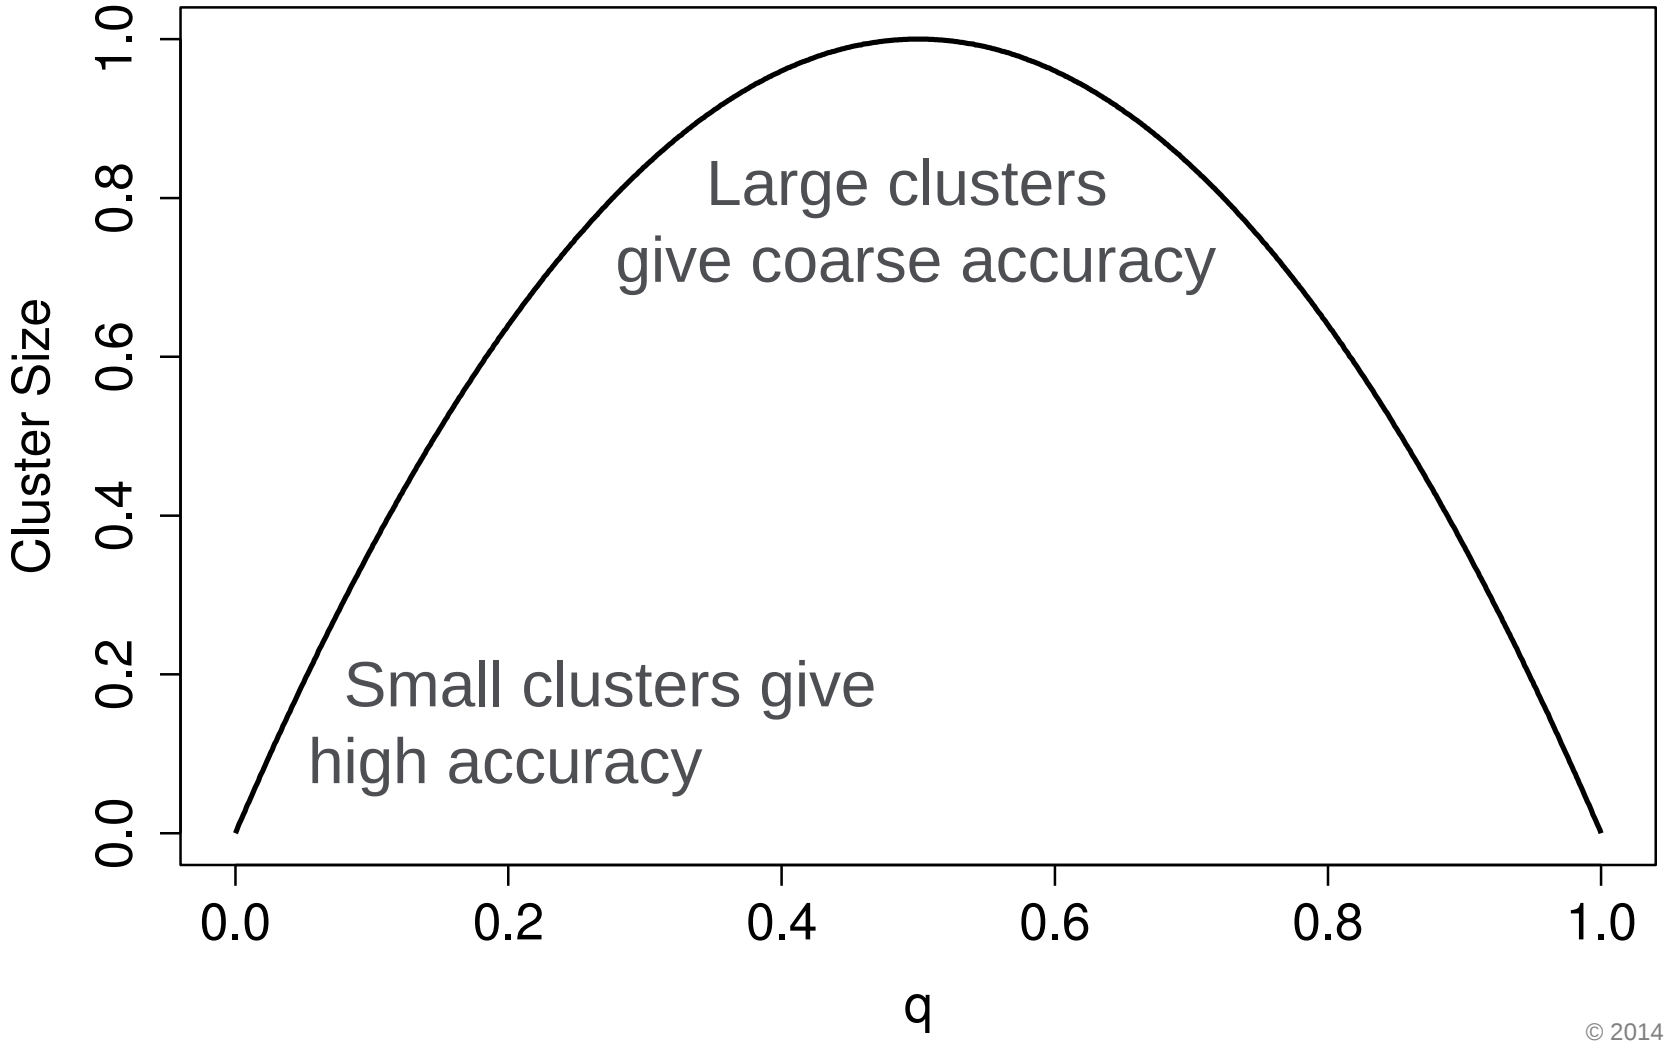

Second-Order Accuracy via Interpolation

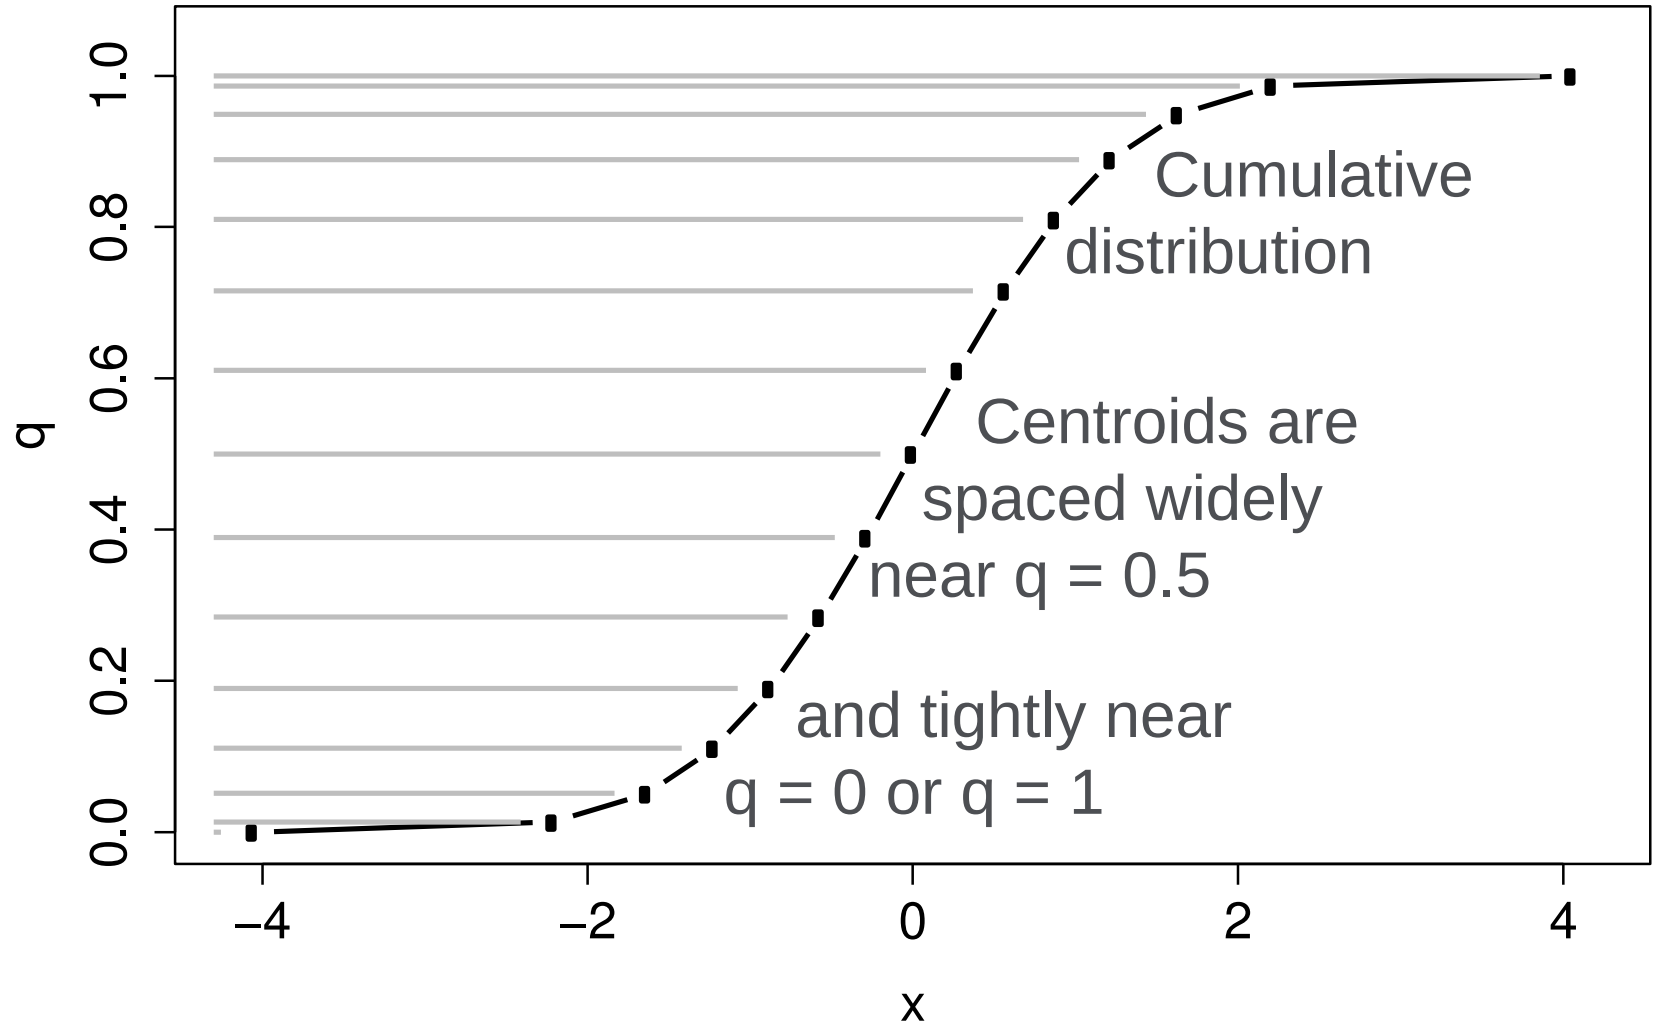

Translation Between Quantile and Cluster

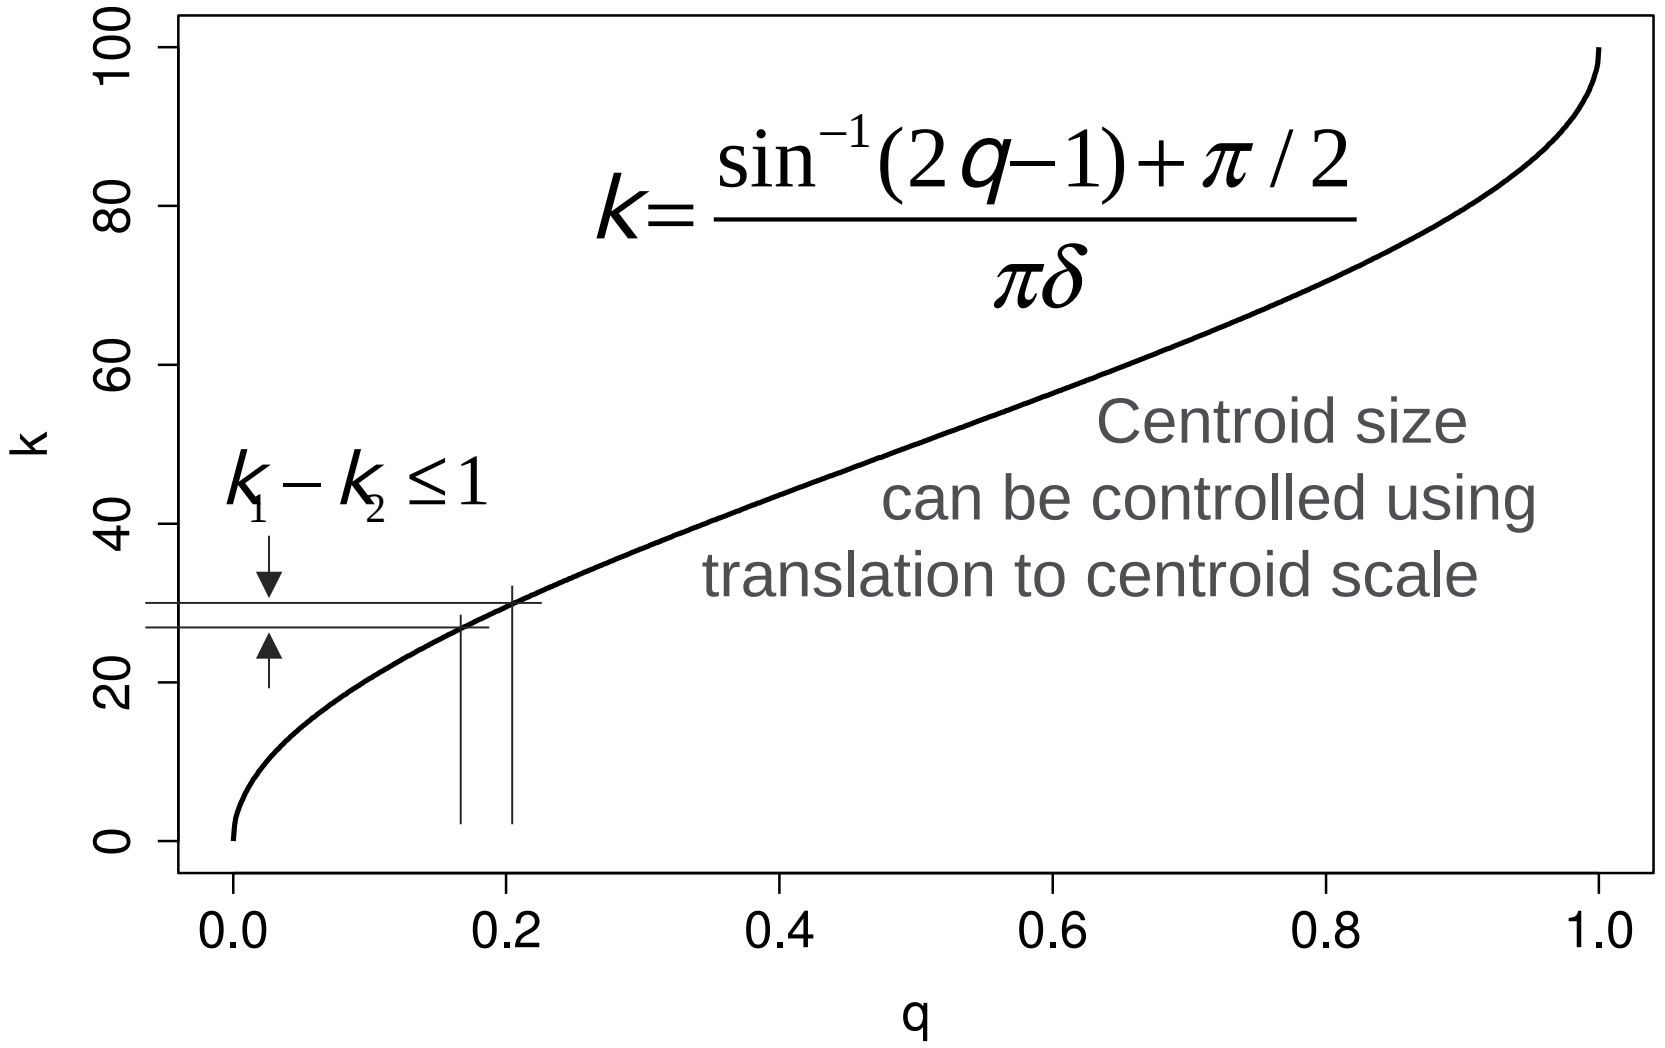

The Algorithm

- Static Buffers
  - n1 new points
  - n2 existing centroids
  - n2 merge space
- Algorithm
  - Collect new points until full
  - Sort new points
  - Merge with existing centroids

    k2 – k1 < 1 criterion for merging
  - Swap centroids and merge space

The Algorithm

- Static Buffers
  - n1 new points
  - n2 existing centroids
  - n2 merge space
- Algorithm
  - Collect new points until full
  - Sort new points
  - Merge with existing centroids

    k2 – k1 < 1 criterion for merging
  - Swap centroids and merge space

- Can be implemented with in-place merge
- Can use approximate q-k mapping for speed
- Completely static memory

Using *t*-digest

Available As

- An aggregator in Elastic Search
- In stream-lib
- As a UDF for Apache Drill (soon!)
- In Apache Mahout
- From Maven Central

```
<dependency>
    <groupId>com.tdunning</groupId>
    <artifactId>t-digest</artifactId>
    <version>3.1</version>
</dependency>
```

The Upshot

- Streaming approximations are important
- Accurate quantiles are important

- The *t*-digest algorithm is simple and very accurate

- You can use it almost anywhere

Special Thanks To

- Otmar Ertl (k2-k1 idea)
- Adrien Grand (best tree implementation)
- Hossman (API improvements)
- Cam Davidson-Pilon (great descriptive blog)

Special Thanks To

- Otmar Ertl (k2-k1 idea)
- Adrien Grand (best tree implementation)
- Hoss (API improvements)
- Cam Davidson-Pilon (great descriptive blog)

- _____ _____ (your name here)

Who I am

Ted Dunning, Chief Applications Architect, MapR Technologies

Email tdunning@mapr.com tdunning@apache.org

Twitter @Ted_Dunning

Apache Mahout https://mahout.apache.org/

Twitter @ApacheMahout

Apache Drill http://incubator.apache.org/drill/

Twitter @ApacheDrill

Q&A

Engage with us!

@mapr   maprtech

mapr-technologies   MapR

tdunning@mapr.com   maprtech