


Big Spatial Data Management


Once upon a
time...



Claudius Ptolemy (AD 90 – AD 168)



Al Idrisi (1099–1165)





TYPVS ORBIS TERRARVM



TERRA AVSTRALIS NONDVM COGNITA

QVID EI POTEST VIDERI MAGNUM IN REBVS HVMANIS, CVI AETERNITAS OMNIS, TOTIVSQUE MVNDI NOTA SIT MAGNITVDO. CICERO:

ARGONAVTICA.



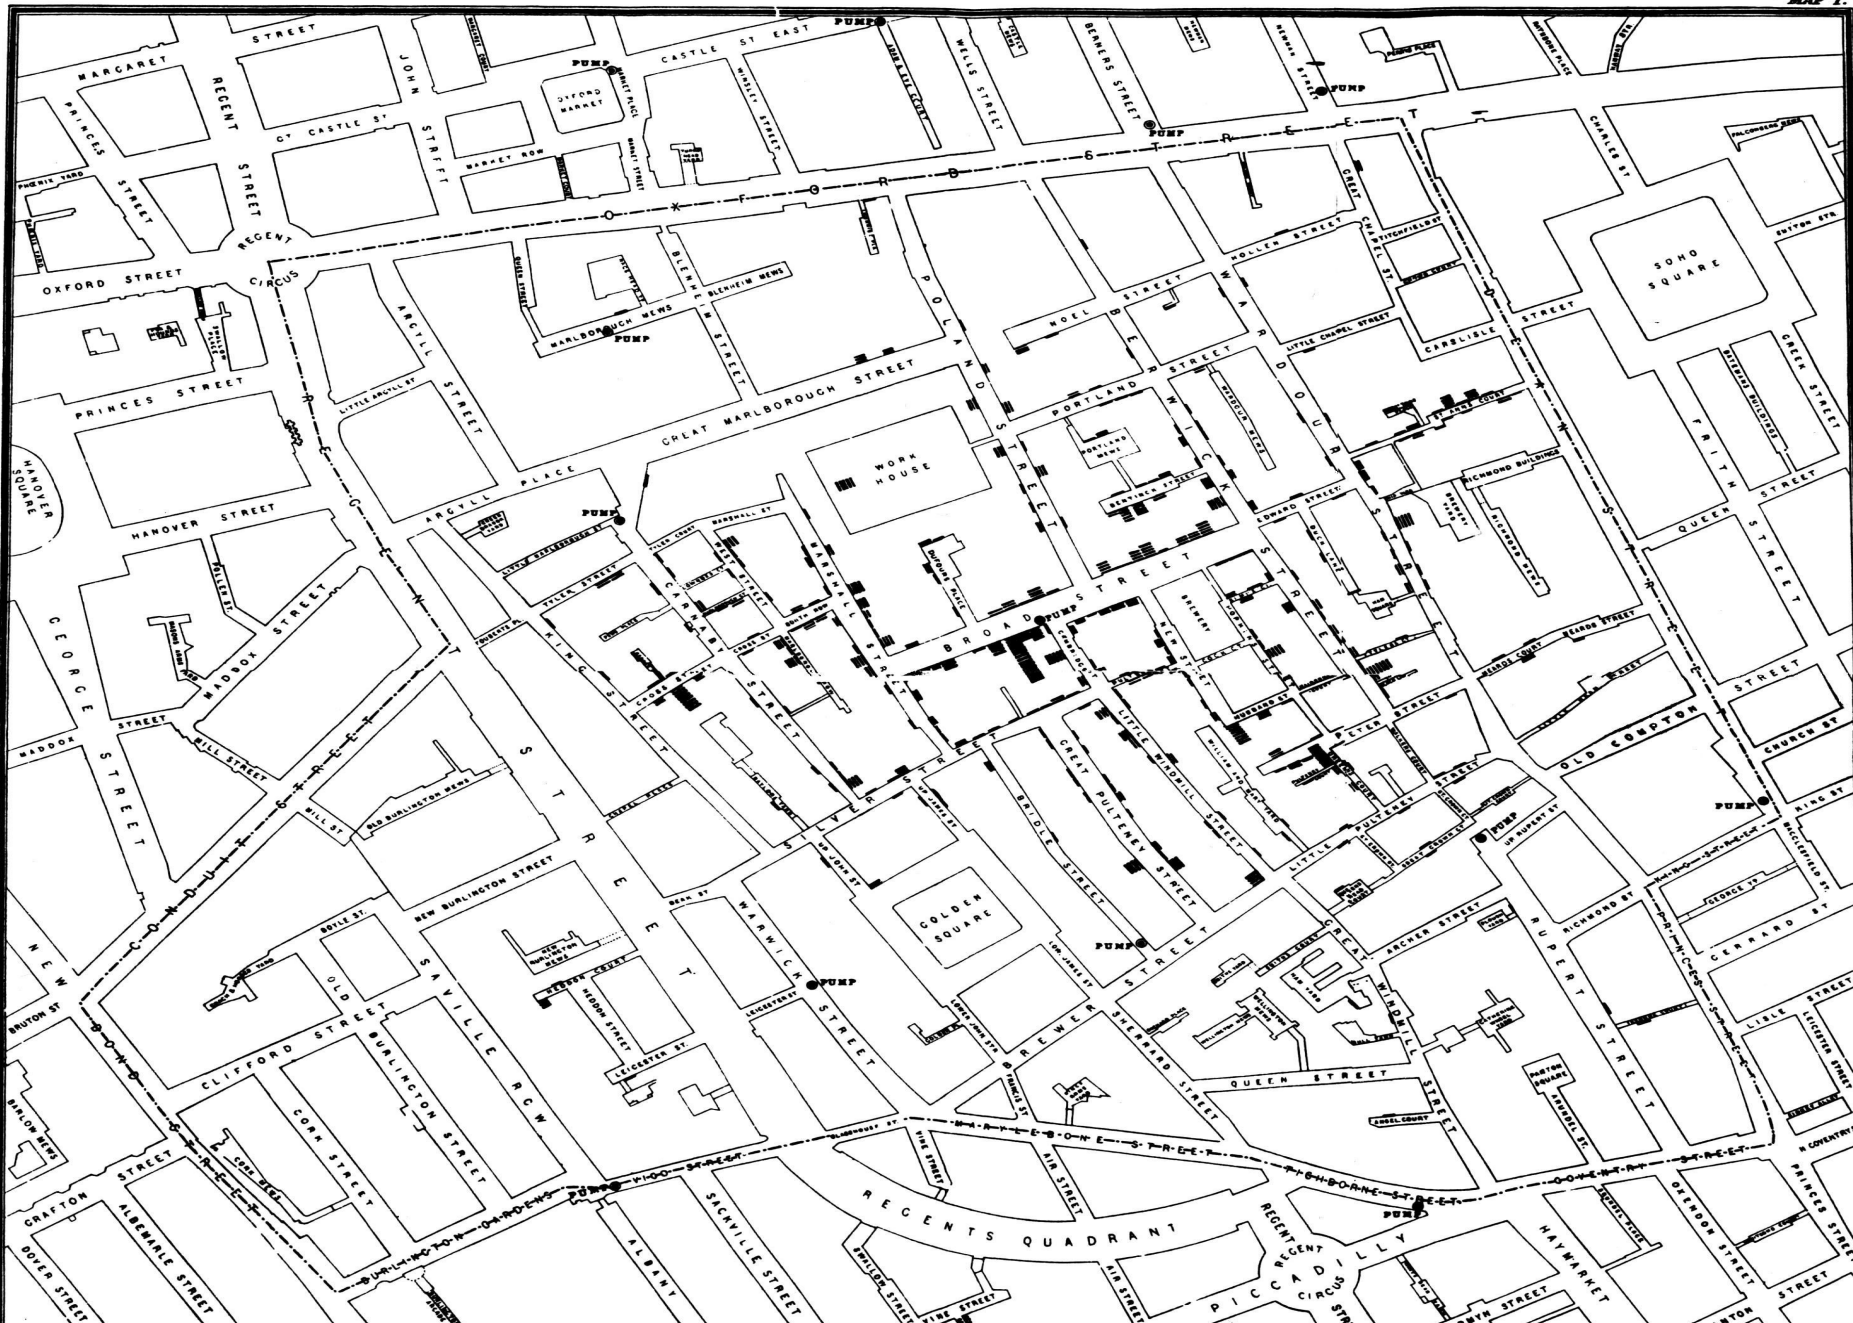
ILLVSTRISSIMO
PRINCIPI CAROLO
COMITI ARENBERGIO,
BARONI SEPTIMONTII,
DOMINO MIRVARTII,
EQVITI AVREI VEL
LERIS, ETC.
ARRAS. OVELINS.
DEWICAR. G. M.

Ex omnibus Geographis Arab. Arab. Arab.

Thracia
Mionium
pelagus

Thraci
Pars
Propontis
Bithynia

Cholera cases in the London epidemic of 1854



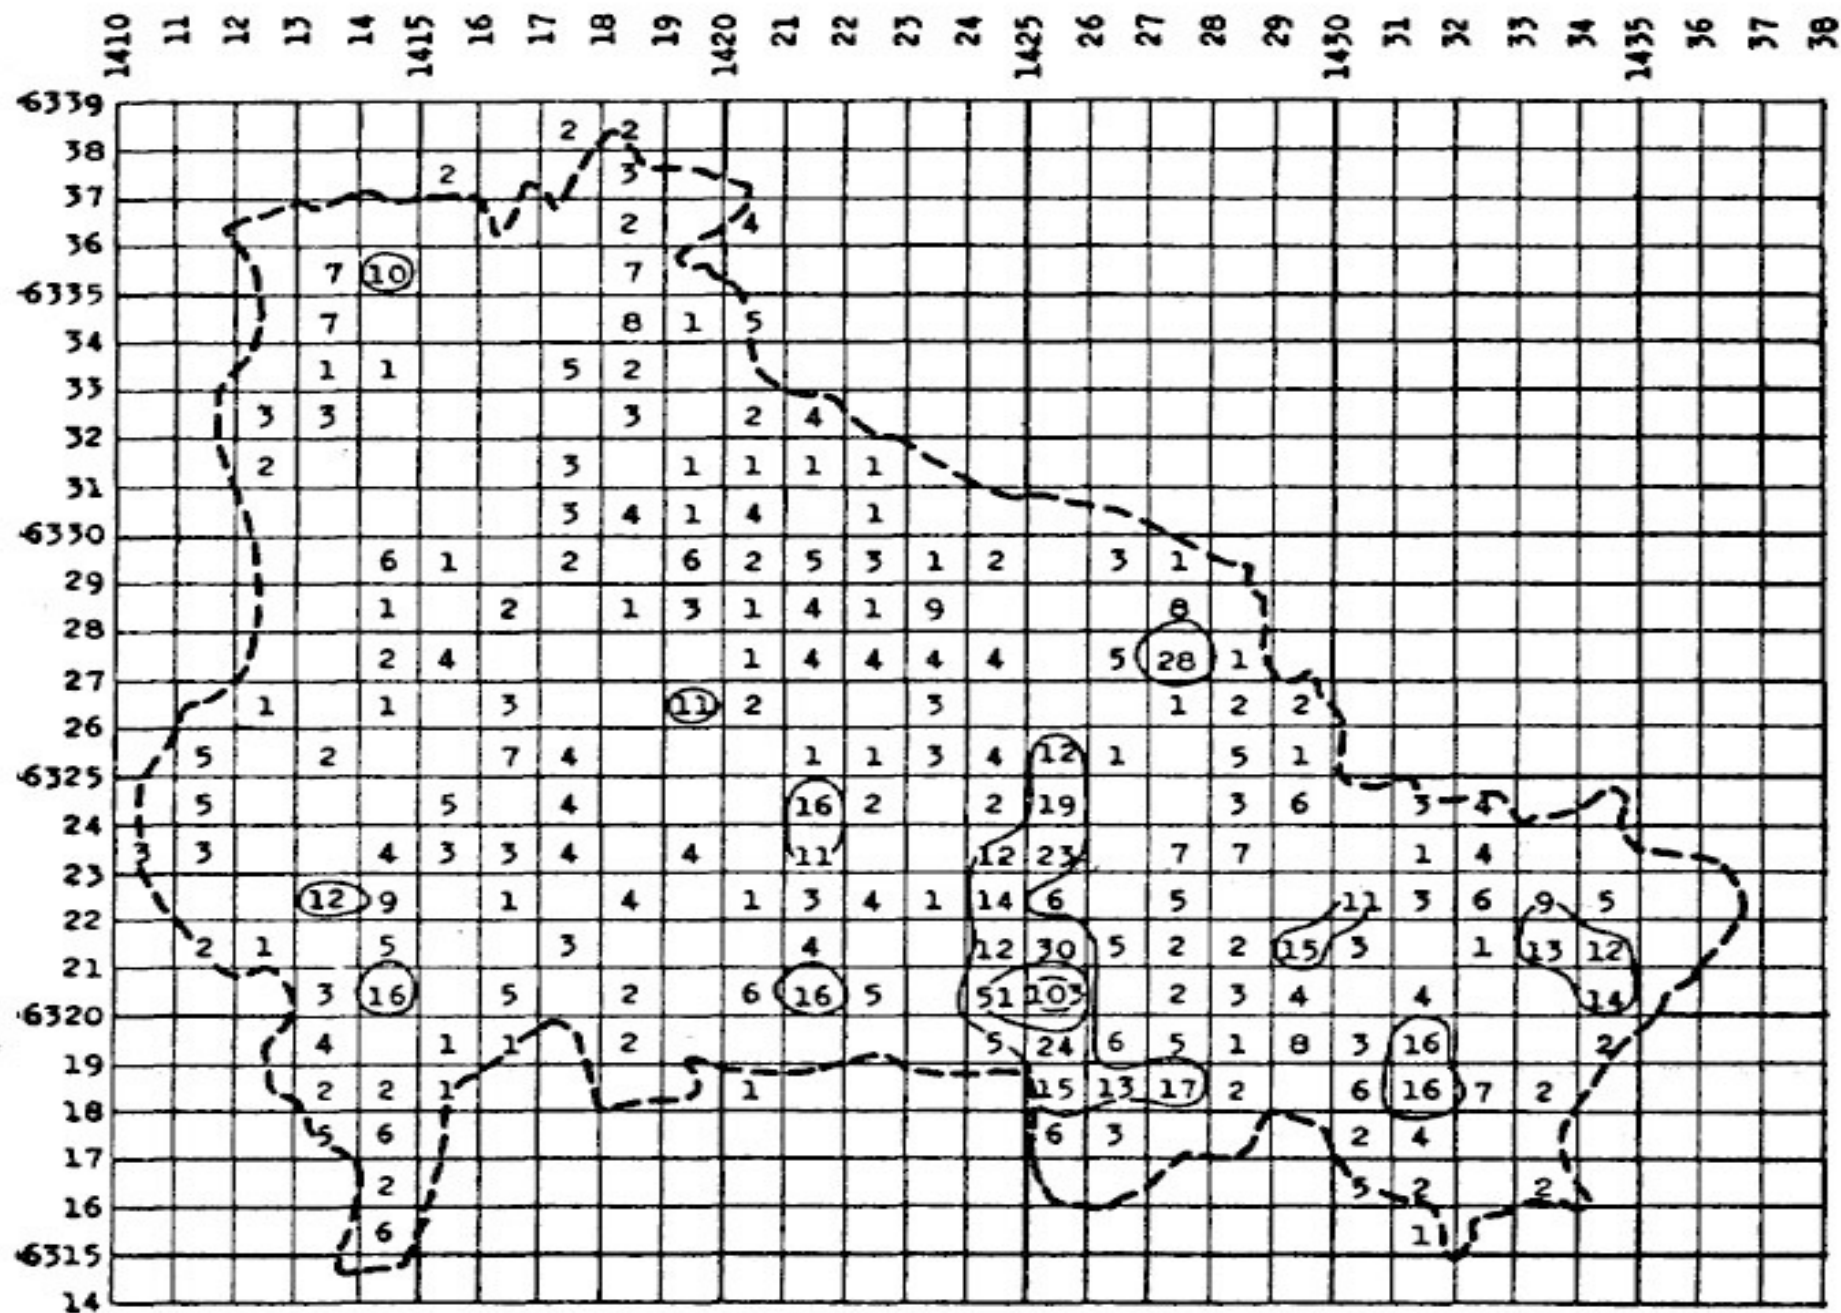


FIGURE 3—Children under 15 years of age in 1940.

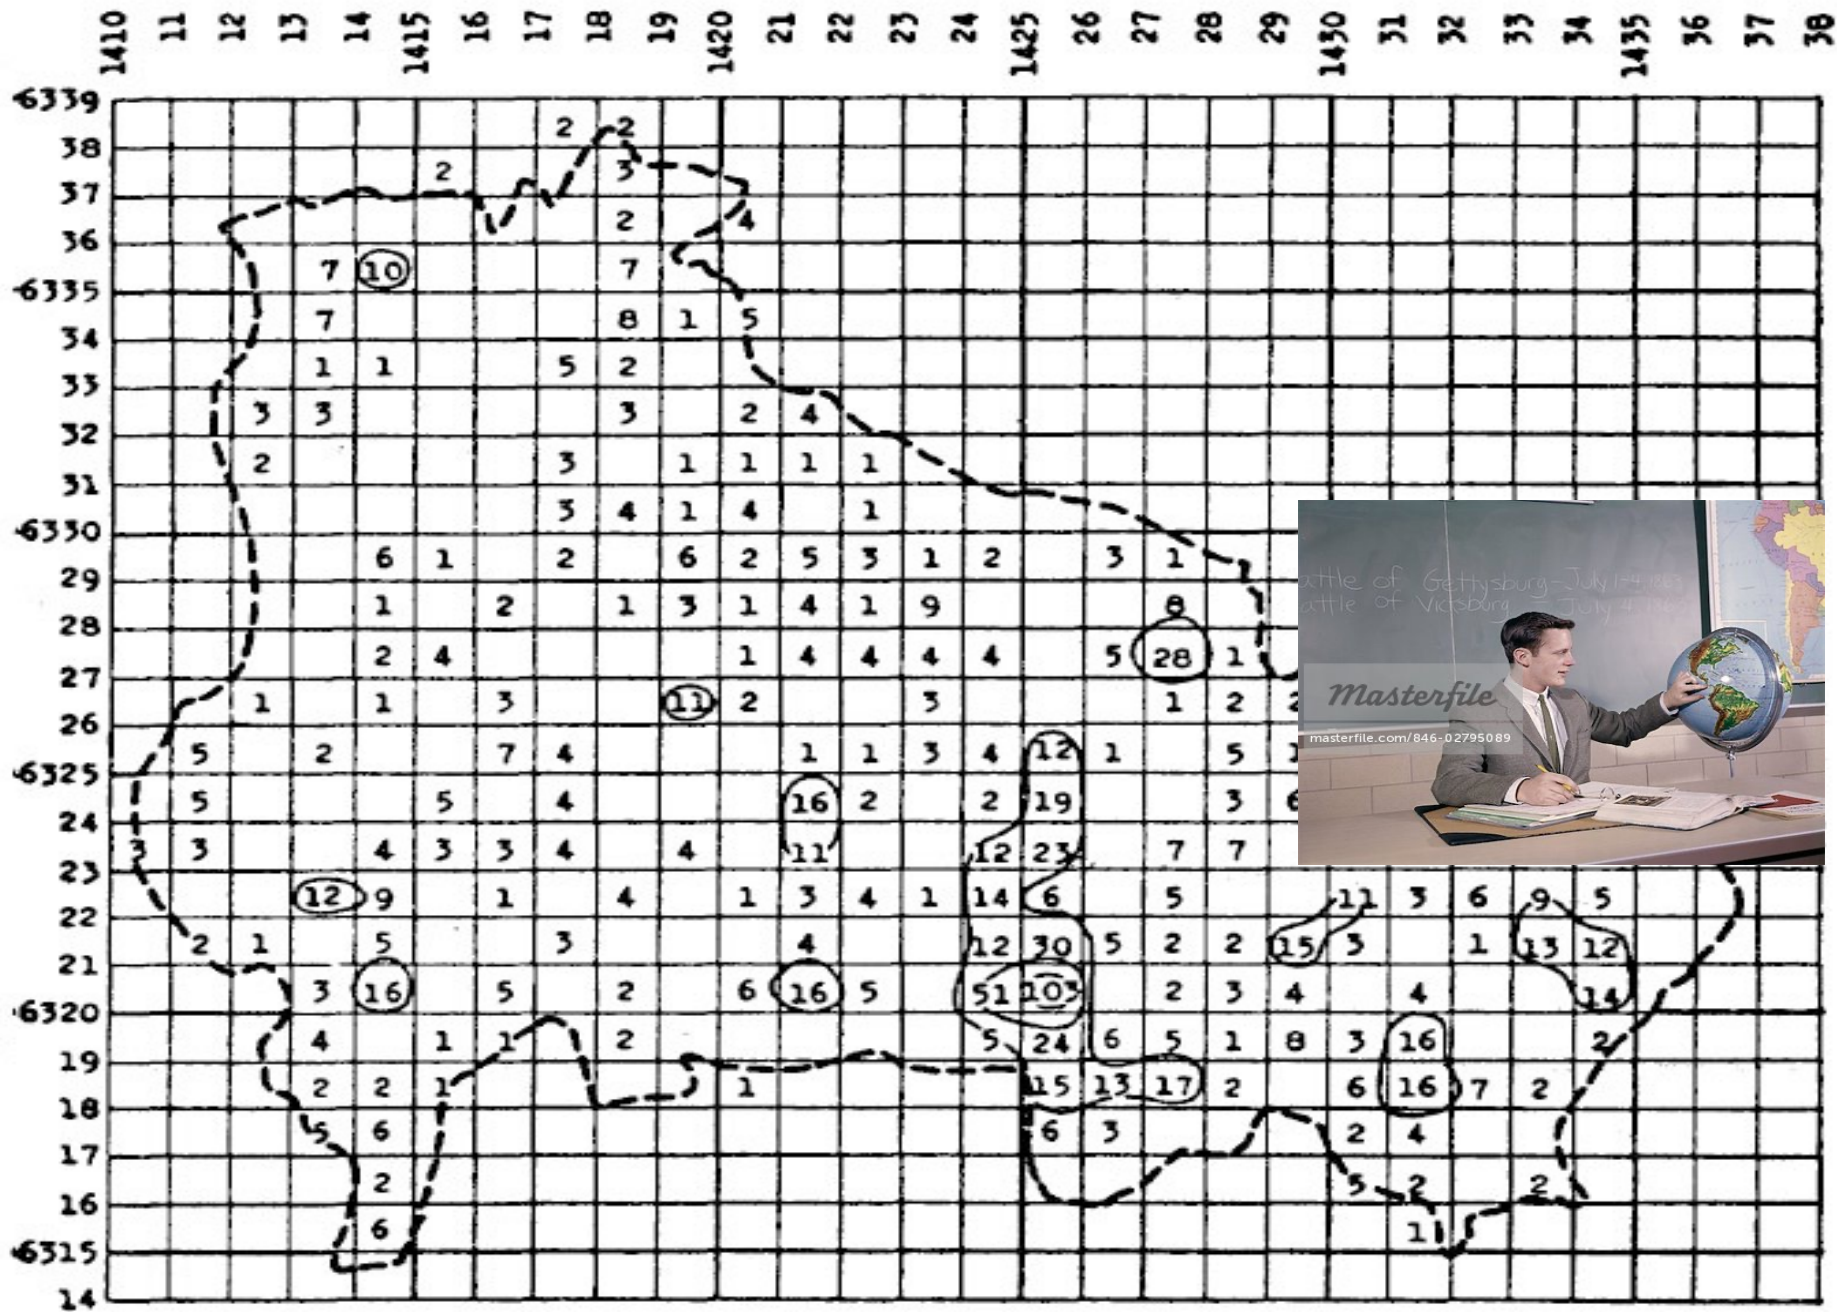
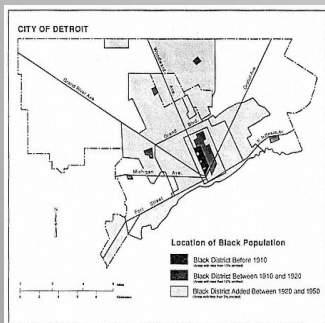

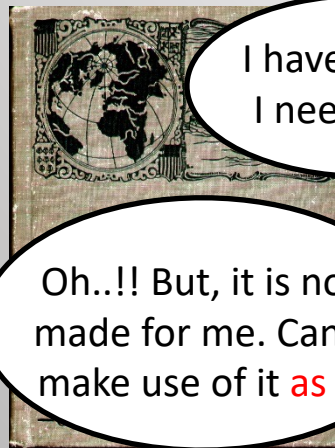


FIGURE 3—Children under 15 years of age in 1940.







Cool **computer** technology..!!
Can I use it in my application?



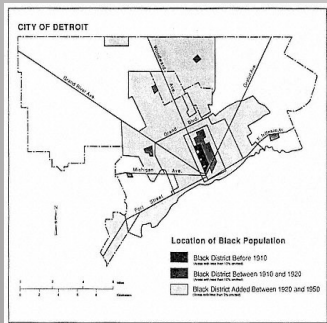
I have **BIG** data.
I need **HELP**..!!



Oh..!! But, it is not made for me. Can't make use of it **as is**



My pleasure. Here it is.






Kindly let me understand your needs

1969

Kindly let me get the technology you have



ESRI



DATABASE MANAGEMENT SYSTEMS



Informix


PostgreSQL





mmm...Let me check with my good friends there.


HELP..!! I have **BIG** data. Your technology is not helping me



Cool **Database** technology..!! Can I use it in my application?



Oh..!! But, it is not made for me. Can't make use of it **as is**



My pleasure. Here it is.



Informix






PostgreSQL






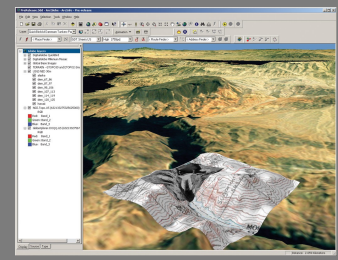















Kindly let me understand your needs

Kindly let me get the technology you have







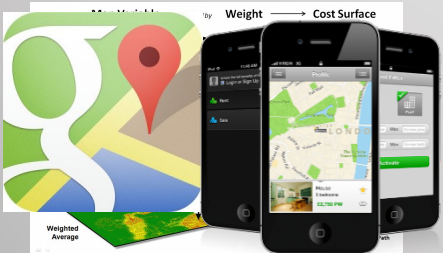








facebook **MapReduce**






Let me check with my **other** good friends there.



Cool **Big Data** technology..!! Can I use it in my application?




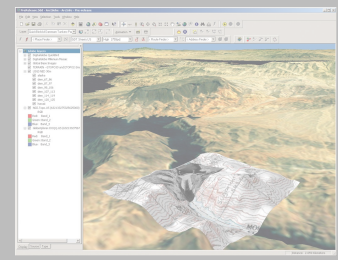




My pleasure. Here it is.

HELP..!! Again, I have **BIG** data. Your technology is not helping me








Sorry, seems like the DBMS technology cannot scale more

Oh..!! But, it's not made for me. Can't make use of it **as is**



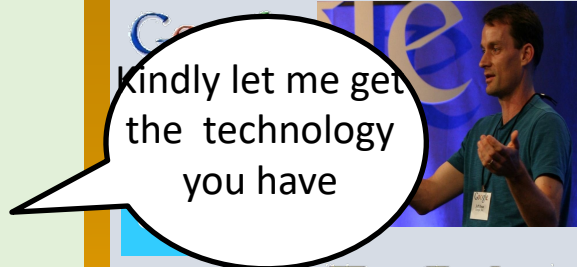


facebook *Map Reduce*











Kindly let me understand your needs





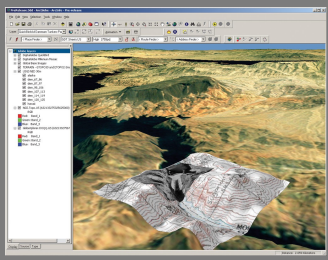



facebook *MapReduce*





Big Spatial Data Management

Google
bing
twitter
facebook *MapReduce*
hadoop
amazon web services™ *HIVE Spark*



Tons of Spatial data out there...

twitter



Geotagged Microblogs




Geotagged Pictures




Medical Data




Smart Phones



Sensor Networks



VGI



Satellite Images



Traffic Data

Spatial Data on Spark

```
val points: RDD[(Double, Double)] = sc.textFile("points.csv")
  .map(l => {
    val coords = l.split(",").map(_.toDouble)
    (coords(0), coords(1))
  })
val xmin, ymin, xmax, ymax: Double = null
val result = points.filter(point => {
  point._1 >= xmin && point._1 < xmax &&
  point._2 >= ymin && point._2 < ymax
})
result.map(pt => s"${pt._1},${pt._2}")
  .saveAsTextFile("output")
```

```
val points: RDD[IFeature] = sc.readCSVPoint("points.csv")
val range = new GeometryFactory().toGeometry(
  new Envelope(xmin, xmax, ymin, ymax))
val results = points.rangeQuery(range)
results.saveAsCSVPoints("output.csv")
```



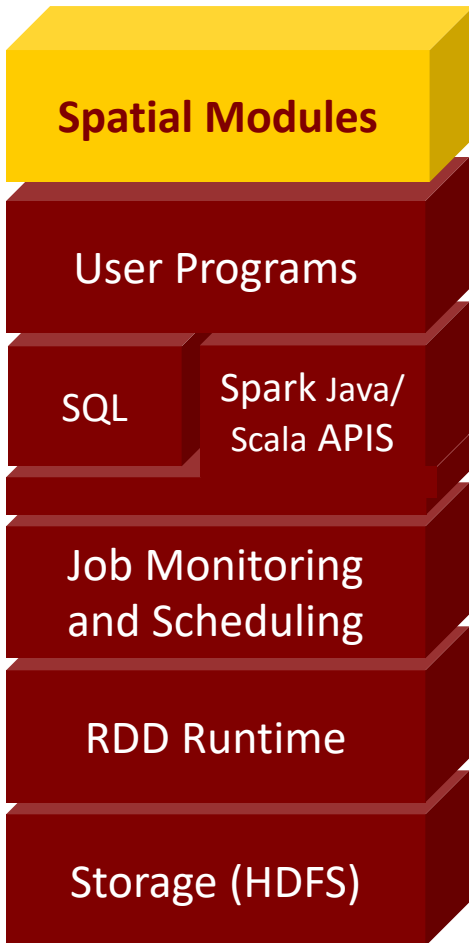
193 seconds



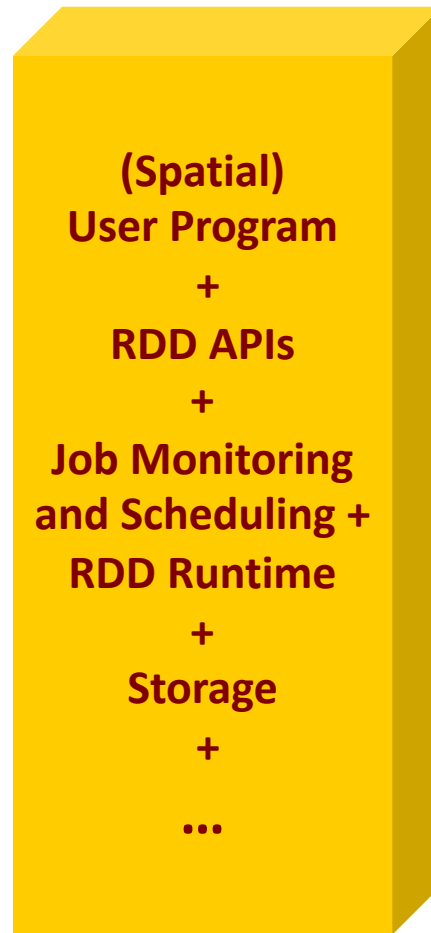
2 seconds

The Built-in Approach of Beast

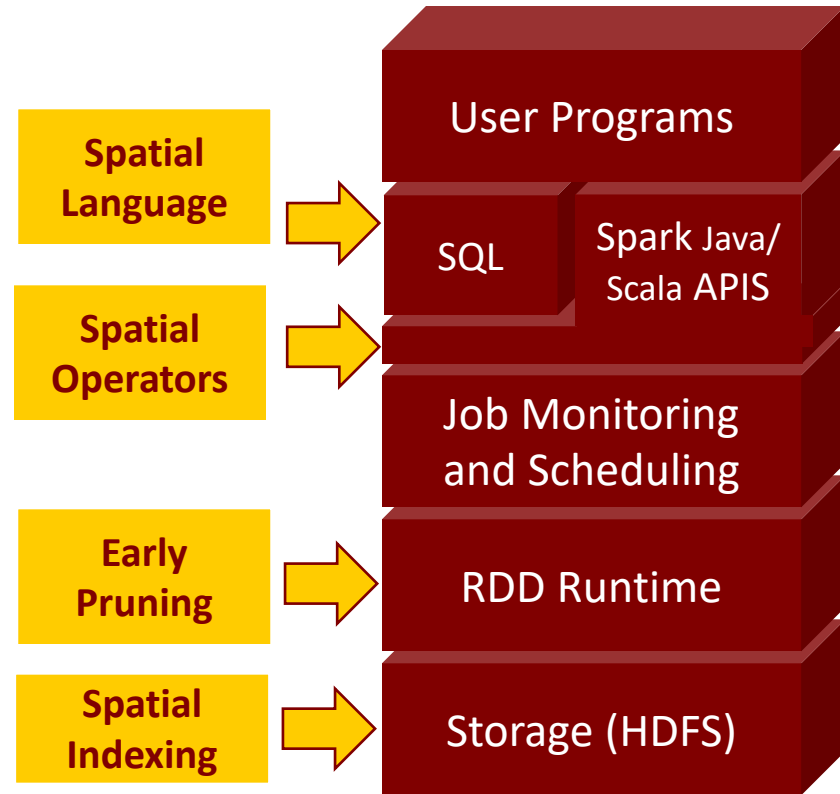
The On-top Approach



From Scratch Approach



The Built-in Approach (Beast)



Domain-specific Big-data

- Spark and similar frameworks are general purpose systems
- They can be customized for a specific domain
- This part is an example of how to customize a big-data system for the domain of spatial data

Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor


Spatial Partitioner & Load Balancer

In-situ Spark Loaders/Writers

Spatial Data Types

Big
Spatial
Data
Apps

Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor

Spatial Partitioner & Load Balancer

In-situ Spark Loaders/Writers

Spatial Data Types

Big
Spatial
Data
Apps

Spatial Data Types

- RDD is flexible enough to allow any user-defined class to be used with RDD
- In Beast, we define the following types
 - Point: n-dimensional point
 - Envelope: n-dimensional box
 - Geometry: Any vector-based geometry
 - Feature: Geometry + attributes



Spatial Data Types



Point



Envelope



Geometry



Feature

Code Samples

```
import org.apache.spark.rdd.RDD
import edu.ucr.cs.bdlab.beast.geolite.IFeature
val buildings: RDD[IFeature] = sc.geojsonFile("buildings.geojson")
```

```
val polygons: SpatialRDD = sc.shapefile("us_counties")
val randomPoints: SpatialRDD = sc.generateSpatialData.
  mbr(polygons.summary).uniform(1000000)
val sjResult = polygons.spatialJoin(randomPoints)
```

Code Samples

```
val counties: SpatialRDD = sc.shapefile("us_counties")
counties.toDataFrame(spark).createOrReplaceTempView("counties")
val counties_areas = spark.sql(
  "SELECT NAME, g, ST_Area(g) FROM counties")
counties_areas.toSpatialRDD.saveAsGeoJSON("us_counties_areas")
```


```
import edu.ucr.cs.bdlab.beast.indexing.RSGrovePartitioner
val partitioned: RDD[(Int, IFeature)] = sc.shapefile("points.shp").
  partitionBy(classOf[RSGrovePartitioner])
```

Code Samples

```
partitioned.saveAsIndex("partitioned_data", "shapefile")  
// To load the data back in another Spark application  
val loadedPartitioned = sc.shapefile("partitioned_data")
```

```
sc.shapefile("us_counties")  
  .plotImage(2000, 2000, "counties.png")
```

Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor

Spatial Partitioner & Load Balancer

In-situ Spark Loaders/Writers

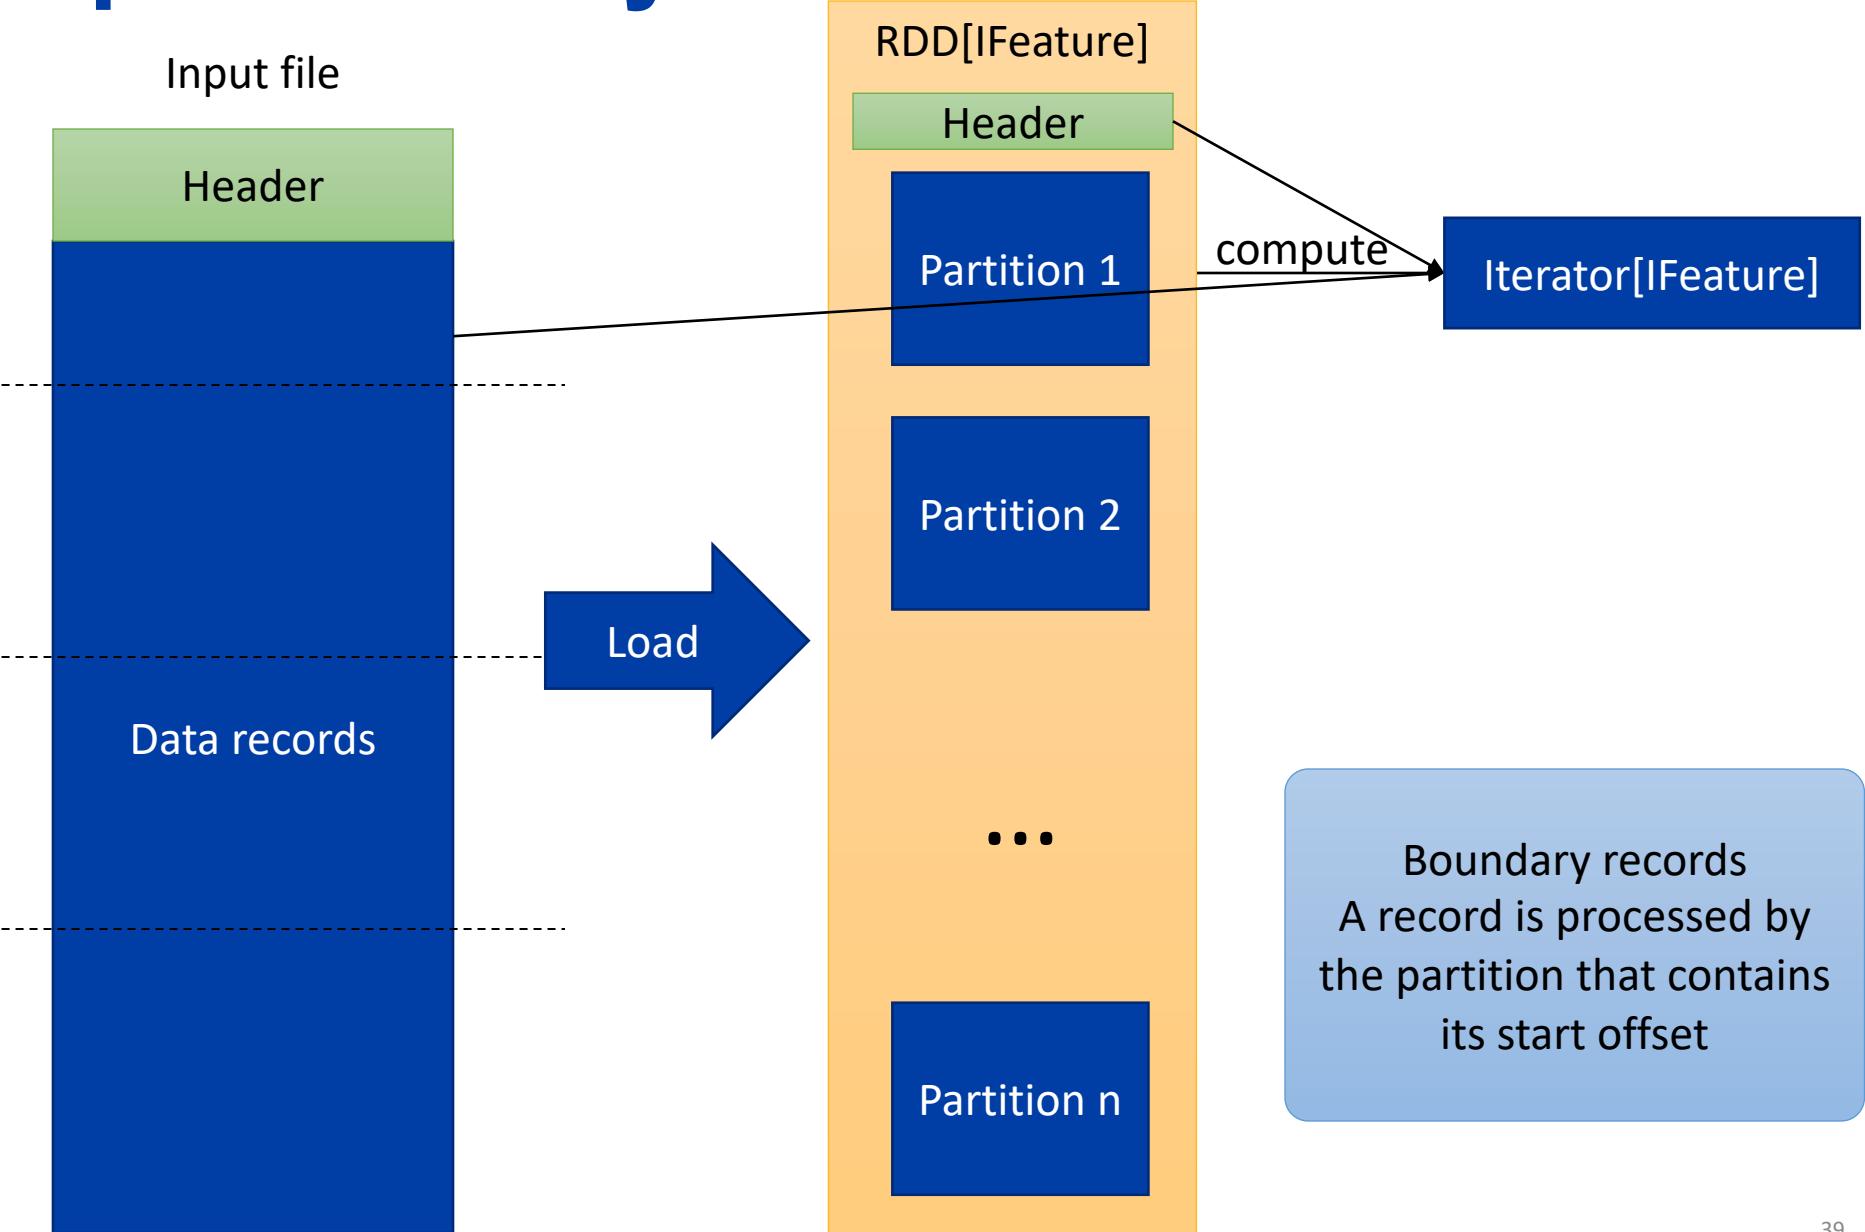
Spatial Data Types

Big
Spatial
Data
Apps

Spark Loaders

- In Spark, a data loader is a top-level RDD that does not depend on any other RDD
- To load data in an input path:
 - Define partitions based on the input metadata
 - Provide a parser for one partition that extracts all records

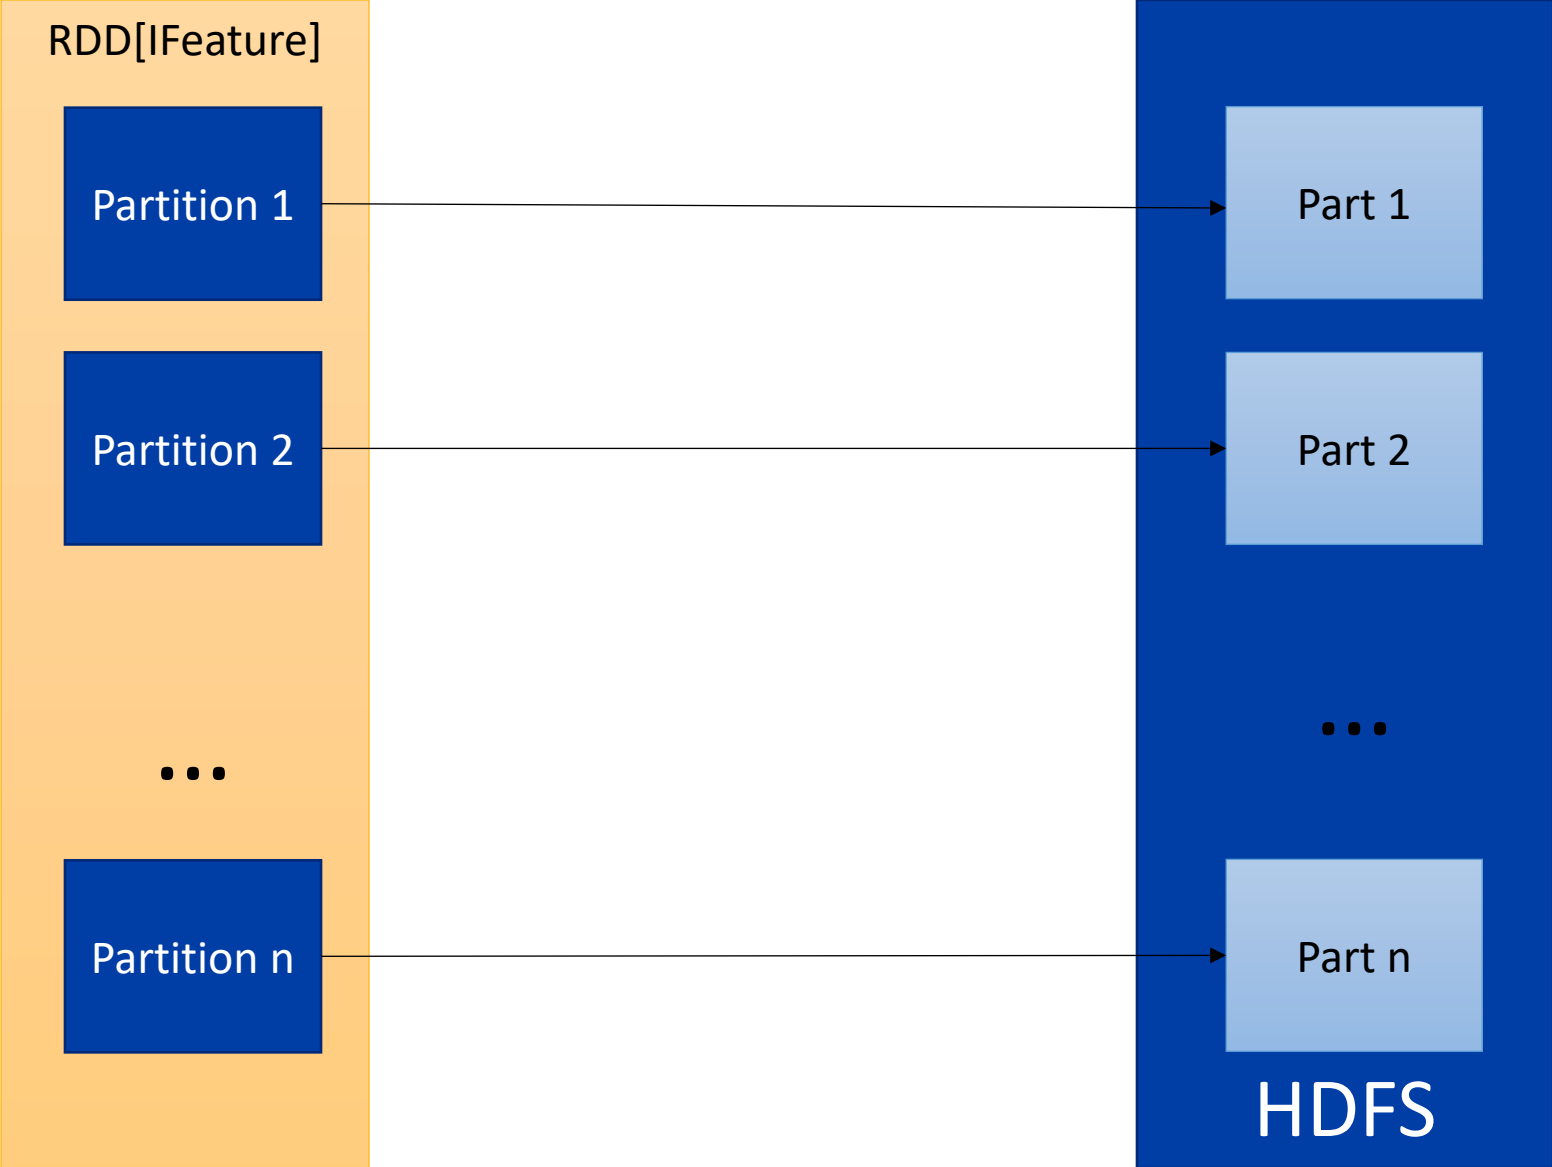
Spatial binary files




Spark Writer

- Implemented as an action
- Operates on RDD[IFeature] and writes all its contents to an output path
- Each partition is written to a separate file

Spark Writer



Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor

Spatial Partitioner & Load Balancer


In-situ Spark Loaders/Writers

Spatial Data Types

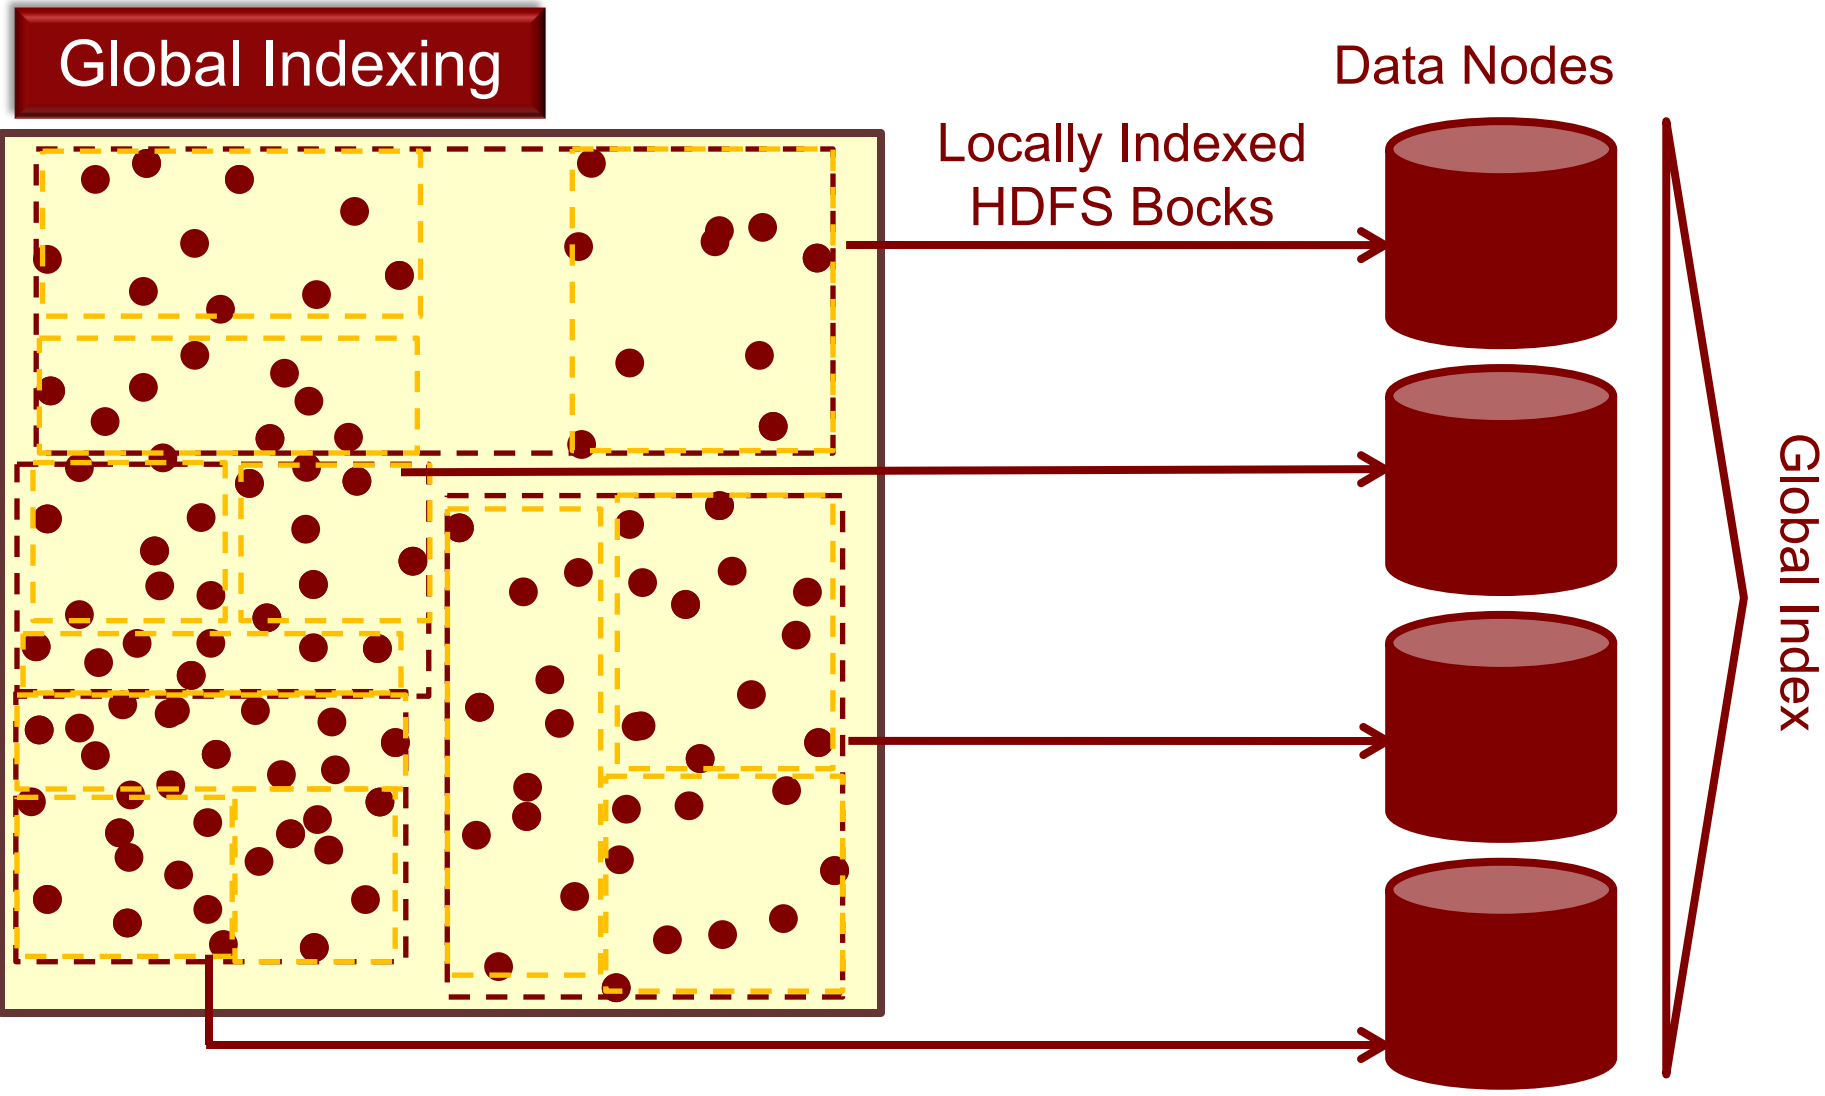
Big
Spatial
Data
Apps

Data Loading in HDFS


- Blindly chops down a big file into 128MB chunks
- Values of records are not considered
- Relevant records are typically assigned to two different blocks
- HDFS is too restrictive where files cannot be modified



Two-layer Index Layout




Uniform Grid



Works only for uniformly distributed data

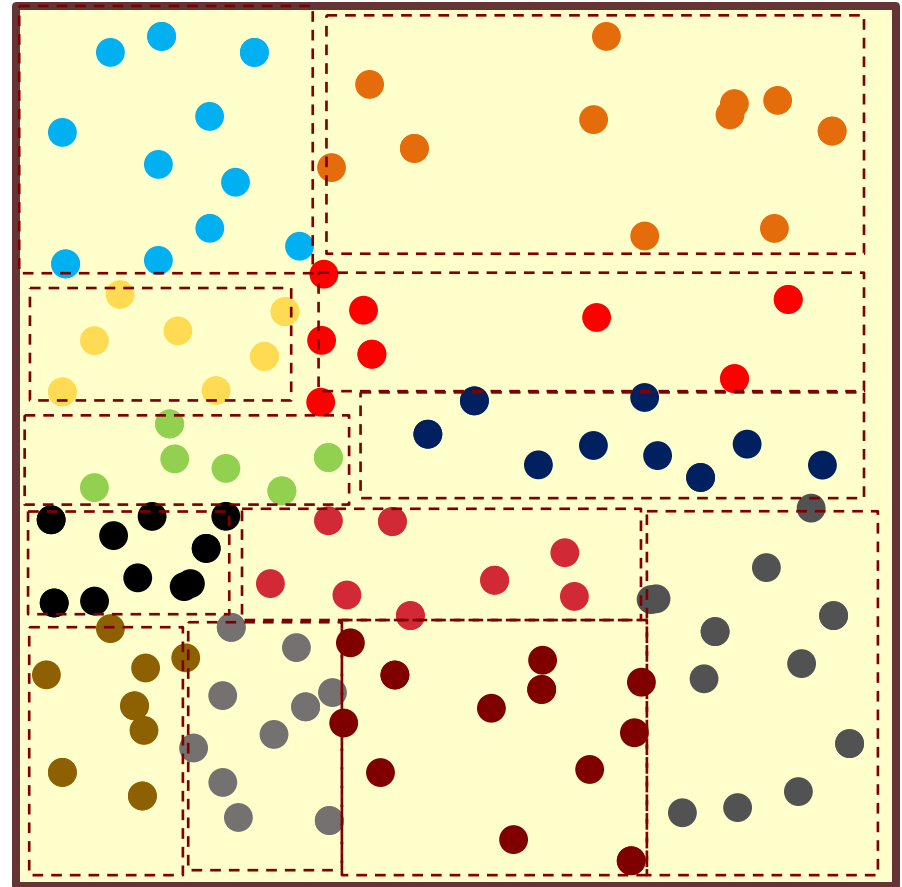
R-tree

- Read a sample
- Partition the sample using an R-tree index
- Use MBR of leaf nodes as partition boundaries for all the data



R-tree


- Read a sample
- Partition the sample using an R-tree index
- Use MBR of leaf nodes as partition boundaries for all the data



R-tree-based Index of a 400 GB road network

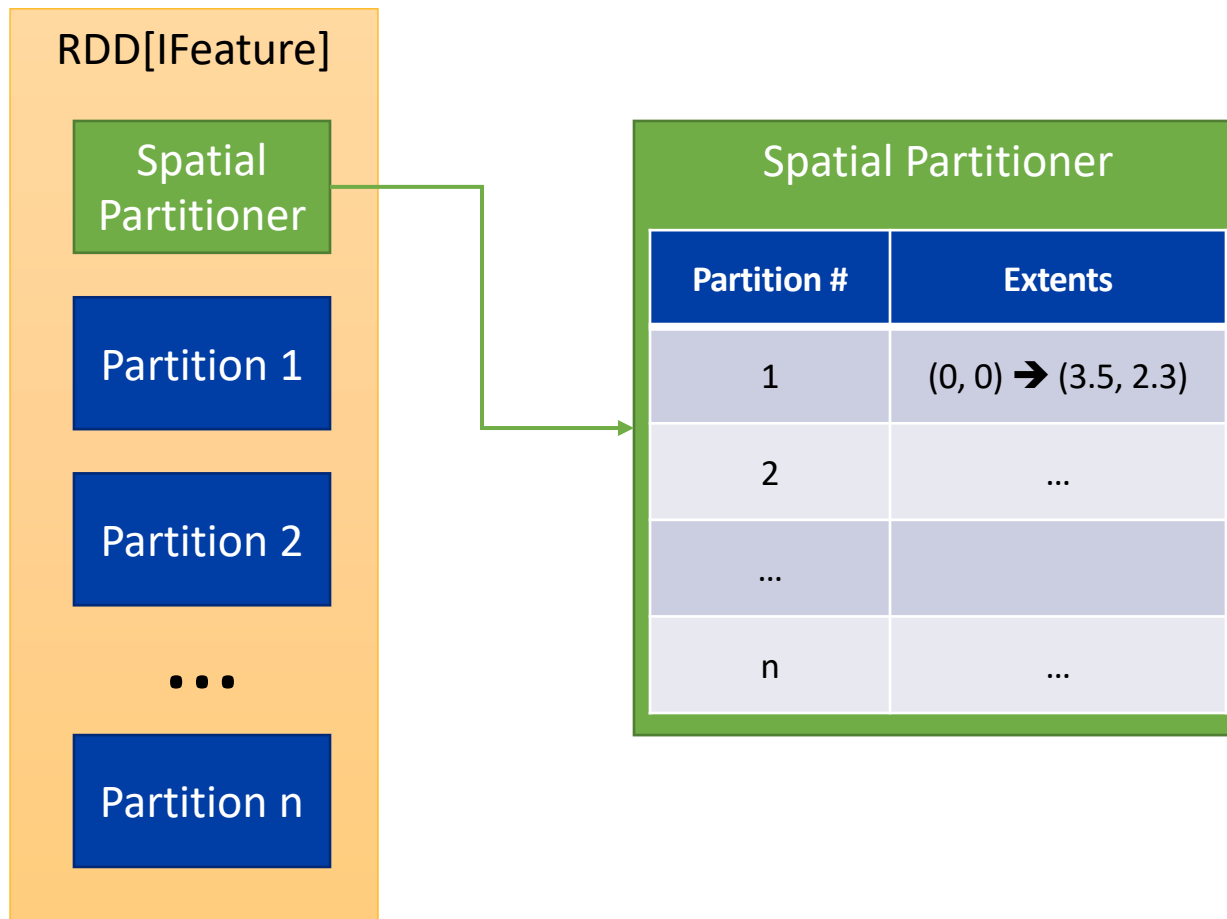


Non-partitioned dataset



Spatially Partitioned RDD

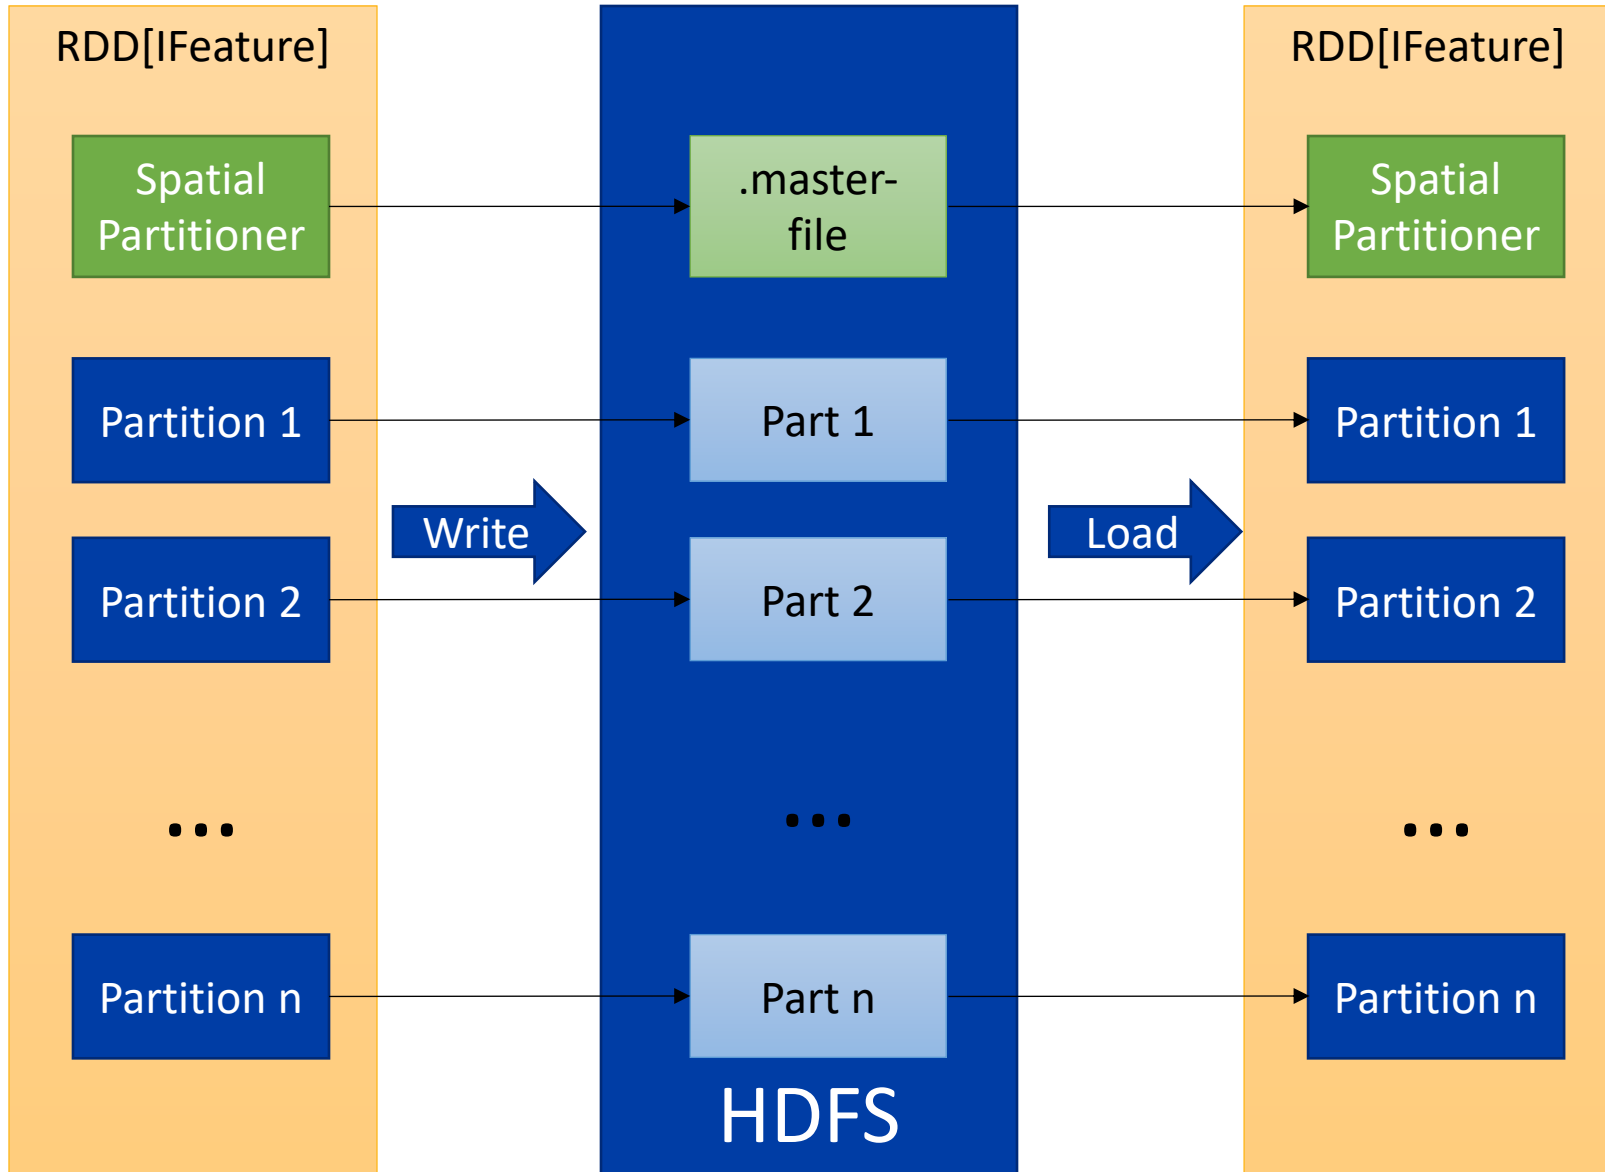
- RDD + Partitioner
- Spark allows custom partitioners




Index Writing and Loading

- Beast provides an option to write an index to disk and read it back
- This gives an option to load an already partitioned RDD

Index Writing and Loading



Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor

Spatial Partitioner & Load Balancer

In-situ Spark Loaders/Writers

Spatial Data Types

Big
Spatial
Data
Apps

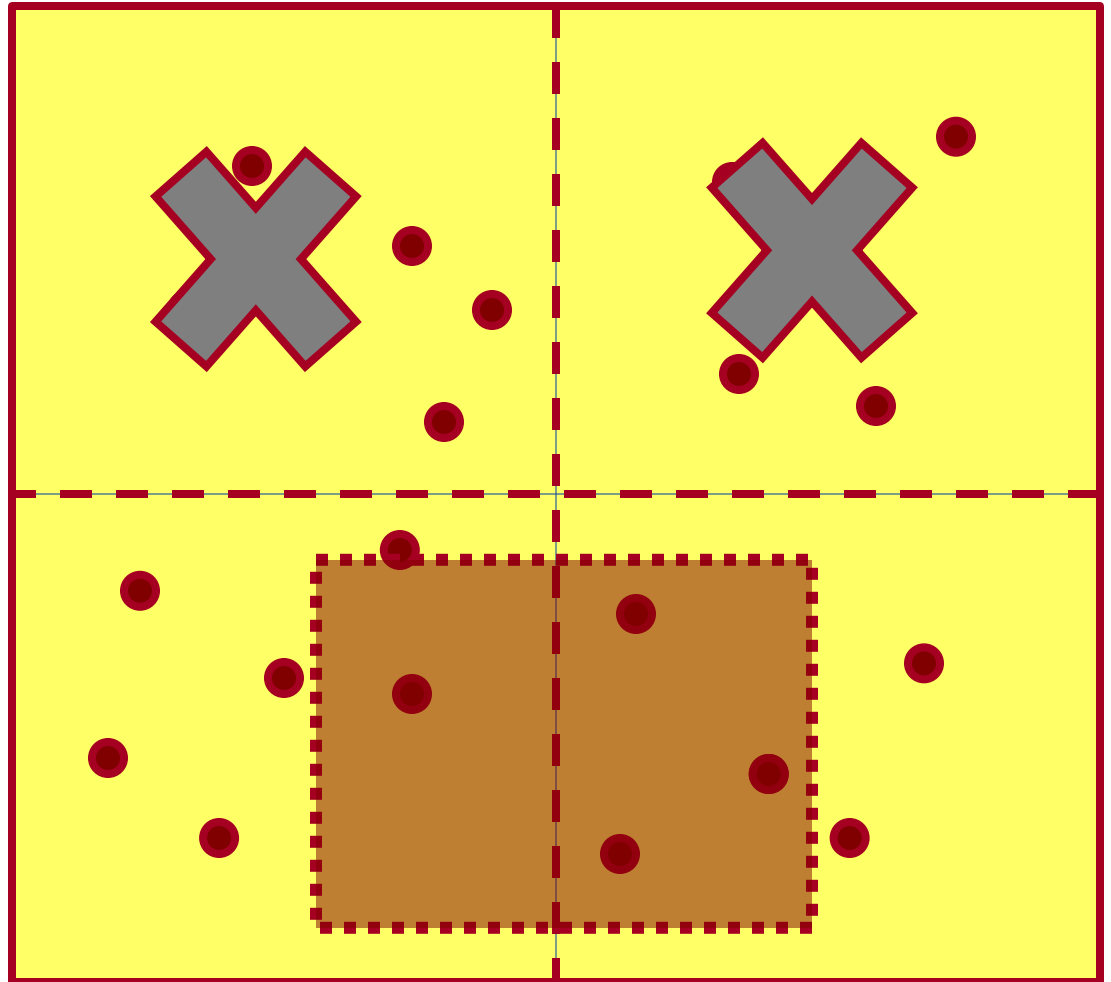
RDD Processing

- Since a spatial RDD is just a regular RDD, all existing transformations and actions can work seamlessly on it
- In addition, we have specialized handling for spatial queries
 - Range Query
 - Spatial Join

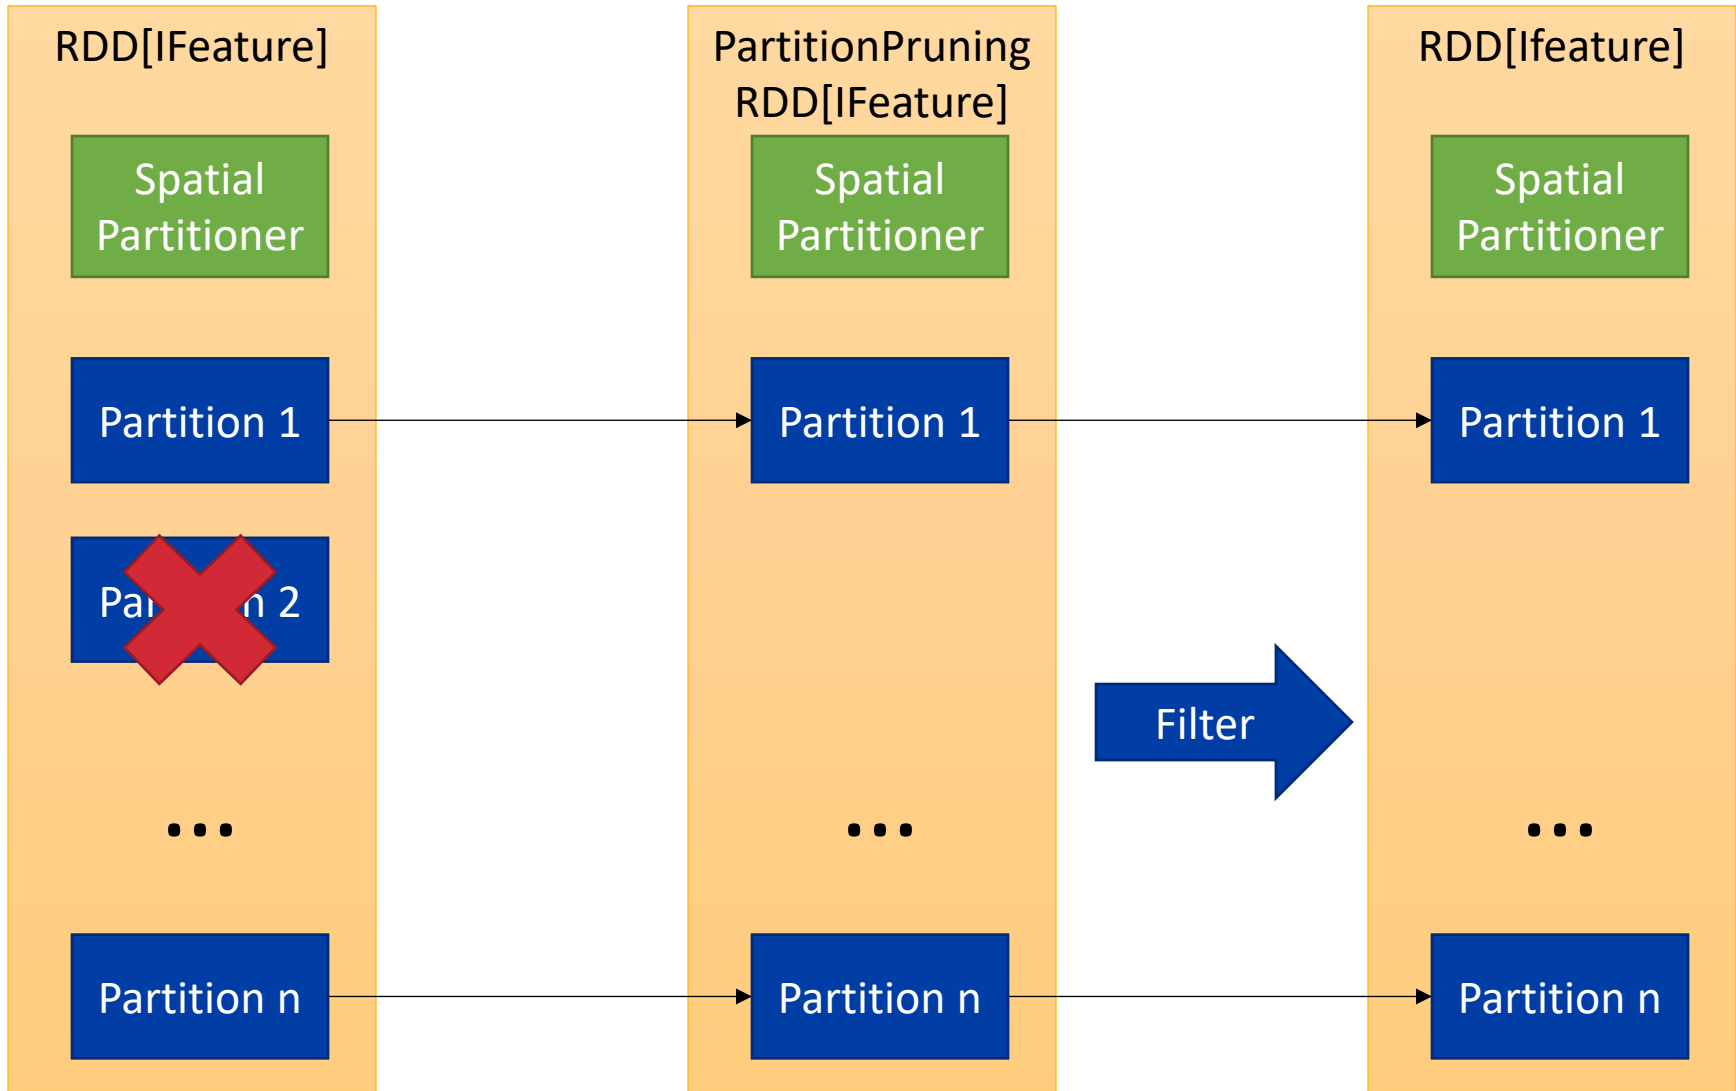
Range Query

Use the **partition information** to prune disjoint partitions

Scan matching partitions in parallel to find matching records



Range Query




Question: Is this narrow or wide dependency?

Spatial Join


Join Directly

Partition – Join




Spatial Join

Join Directly



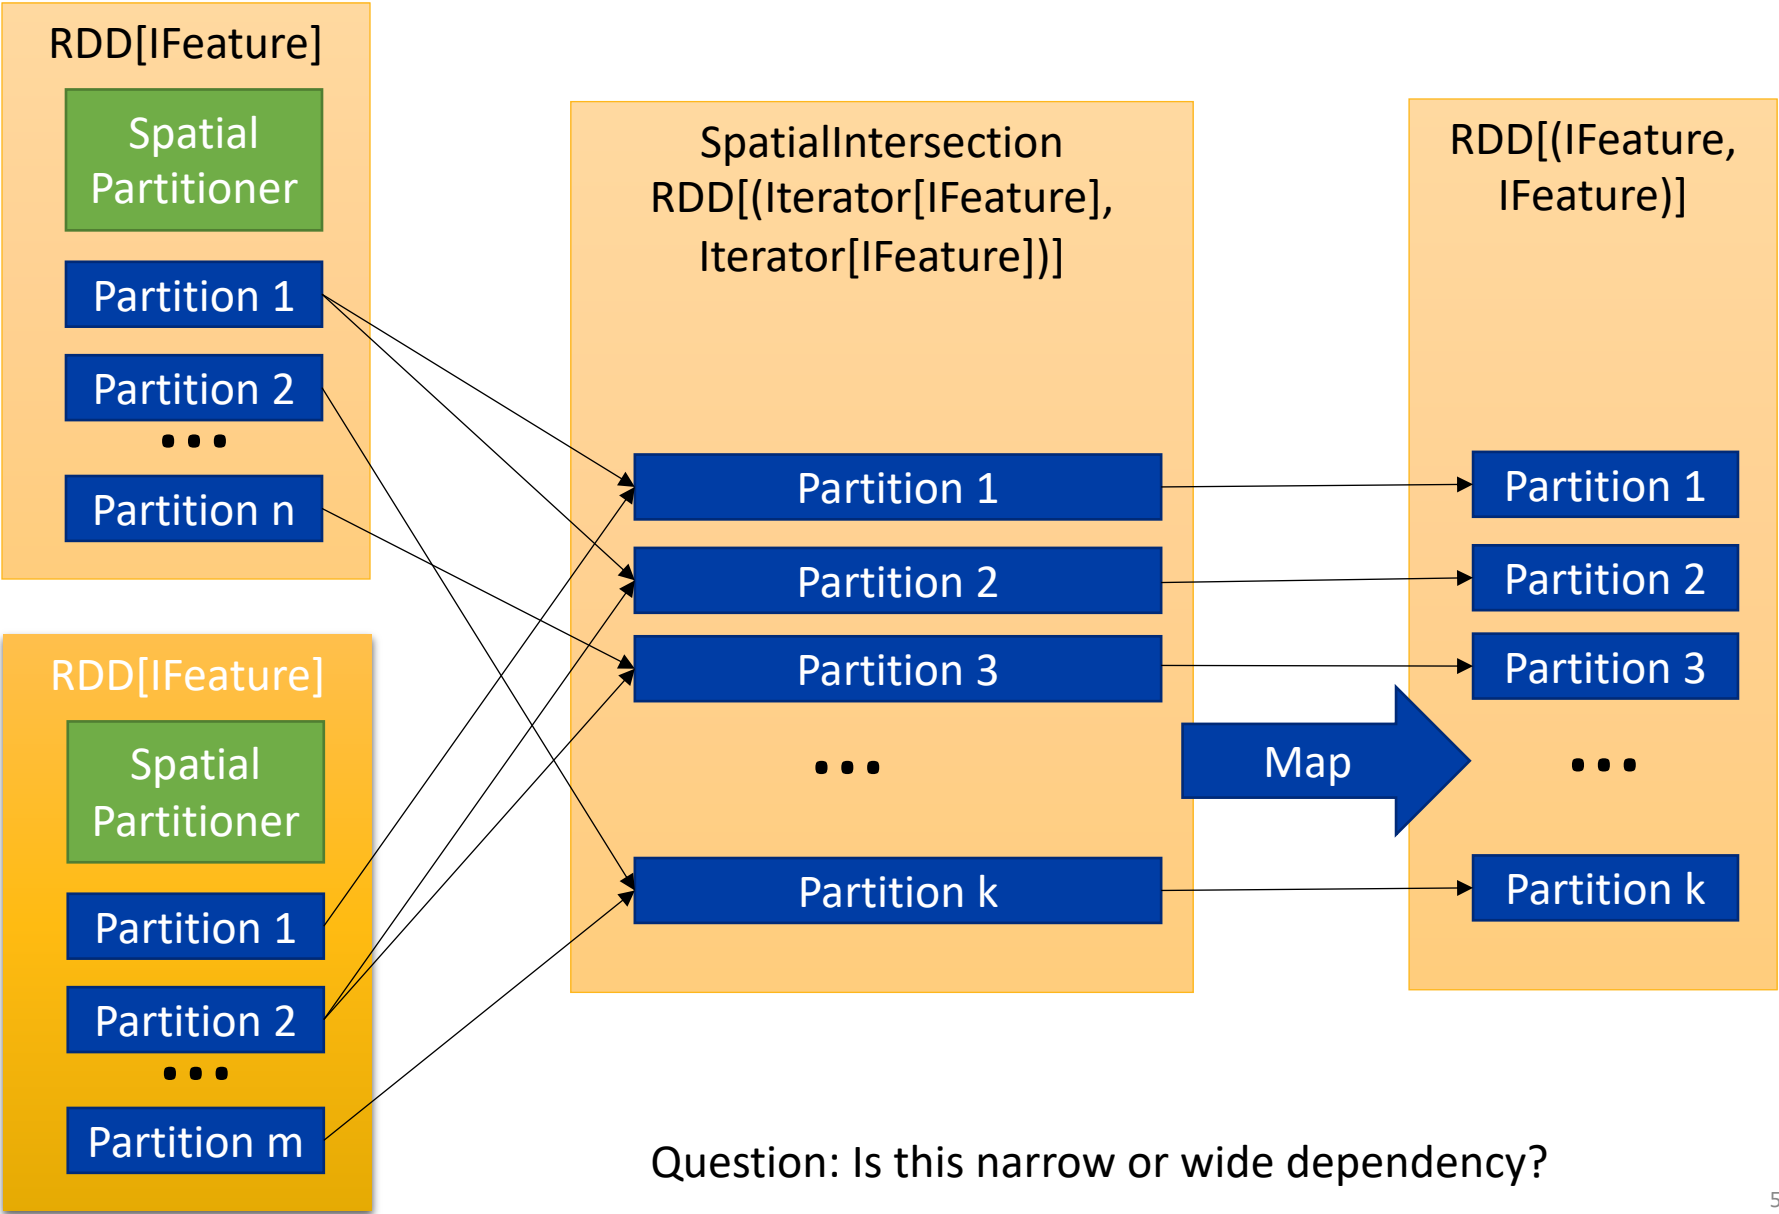
Total of 36 overlapping pairs

Partition – Join

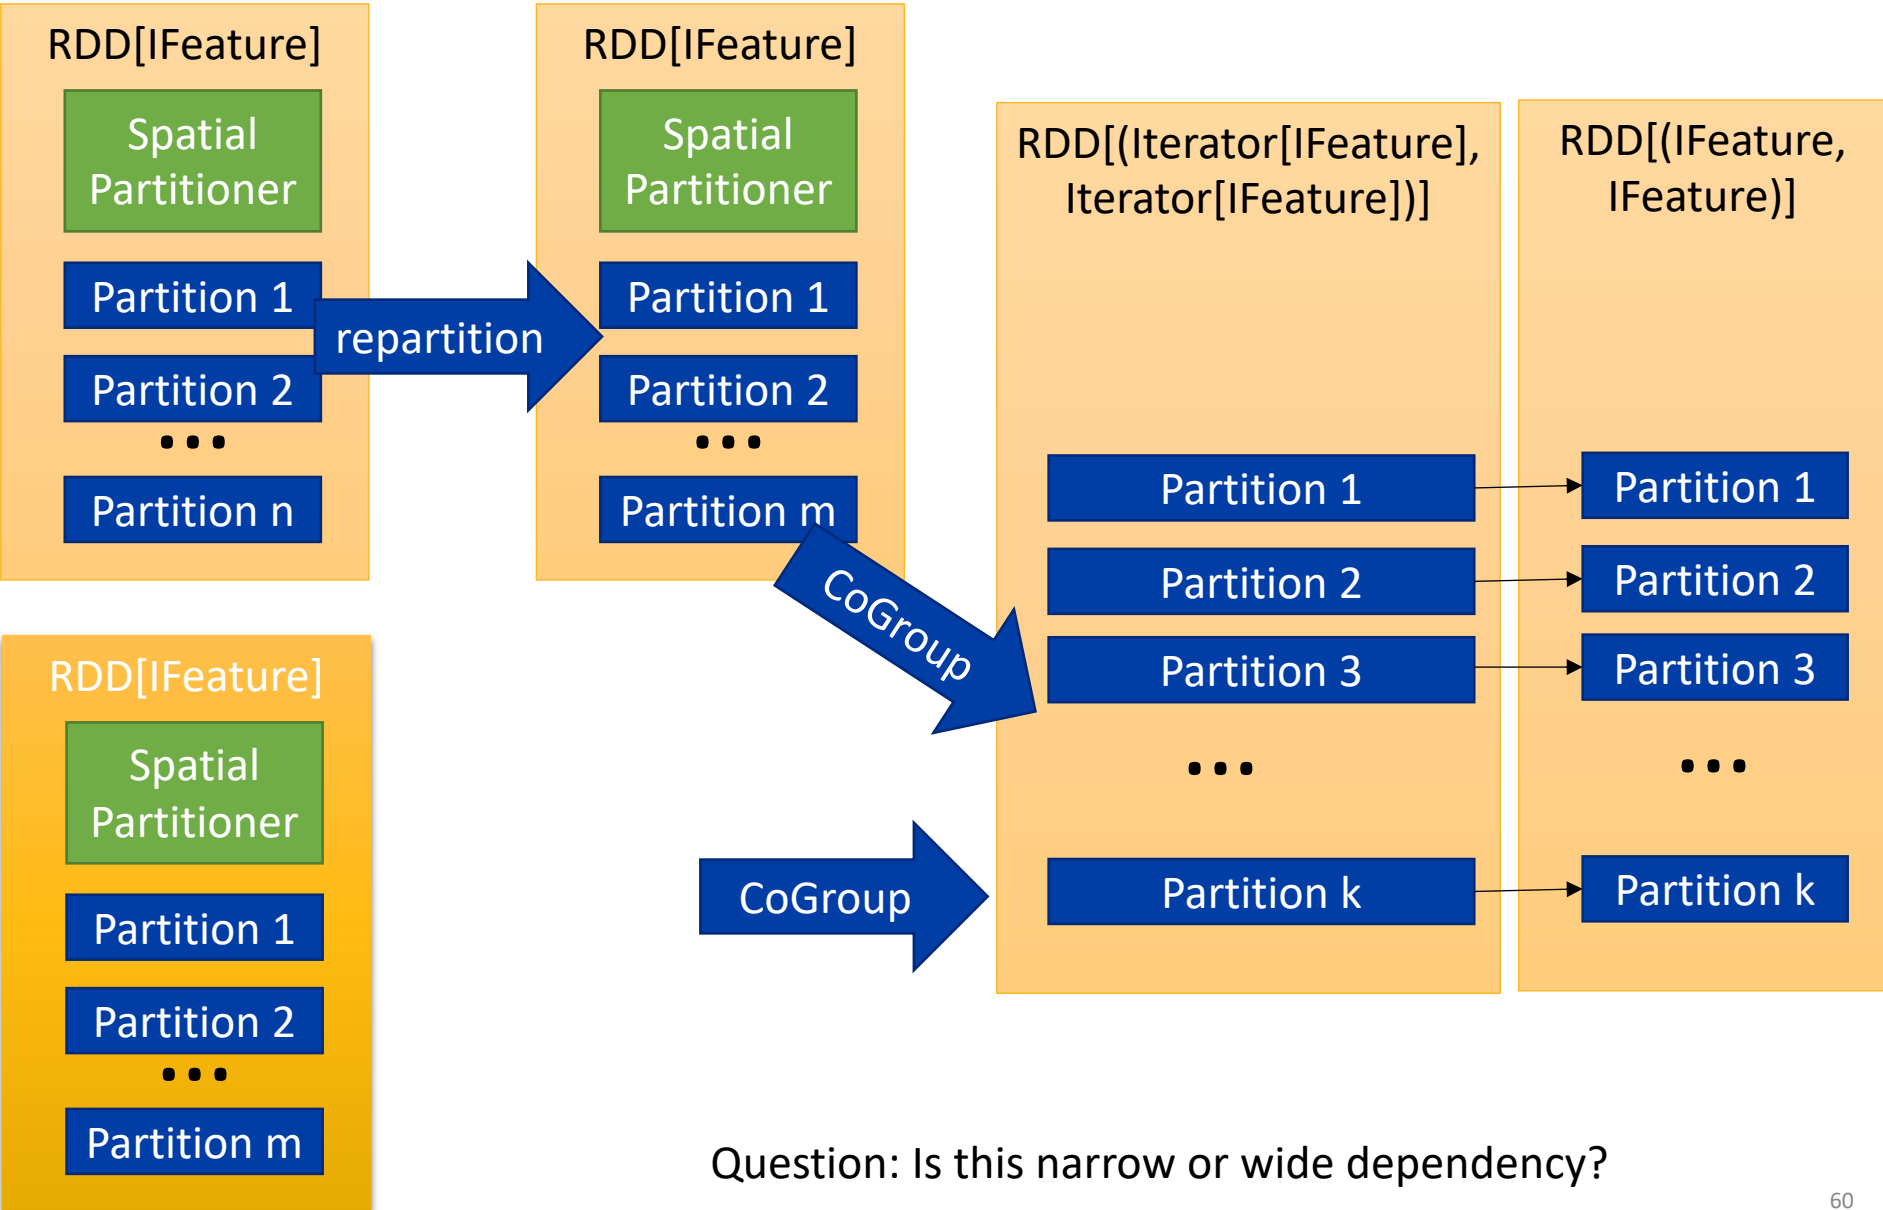


Only 16 overlapping pairs

Join Directly




Join Directly



Question: Is this narrow or wide dependency?

Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor




Spatial Partitioner & Load Balancer

In-situ Spark Loaders/Writers

Spatial Data Types

Big
Spatial
Data
Apps


Visualization in HadoopViz




The goal of **HadoopViz** is not to propose **new visualization** techniques, instead its goal is to **scale out** existing techniques.



Scatter Plot



Vector Map

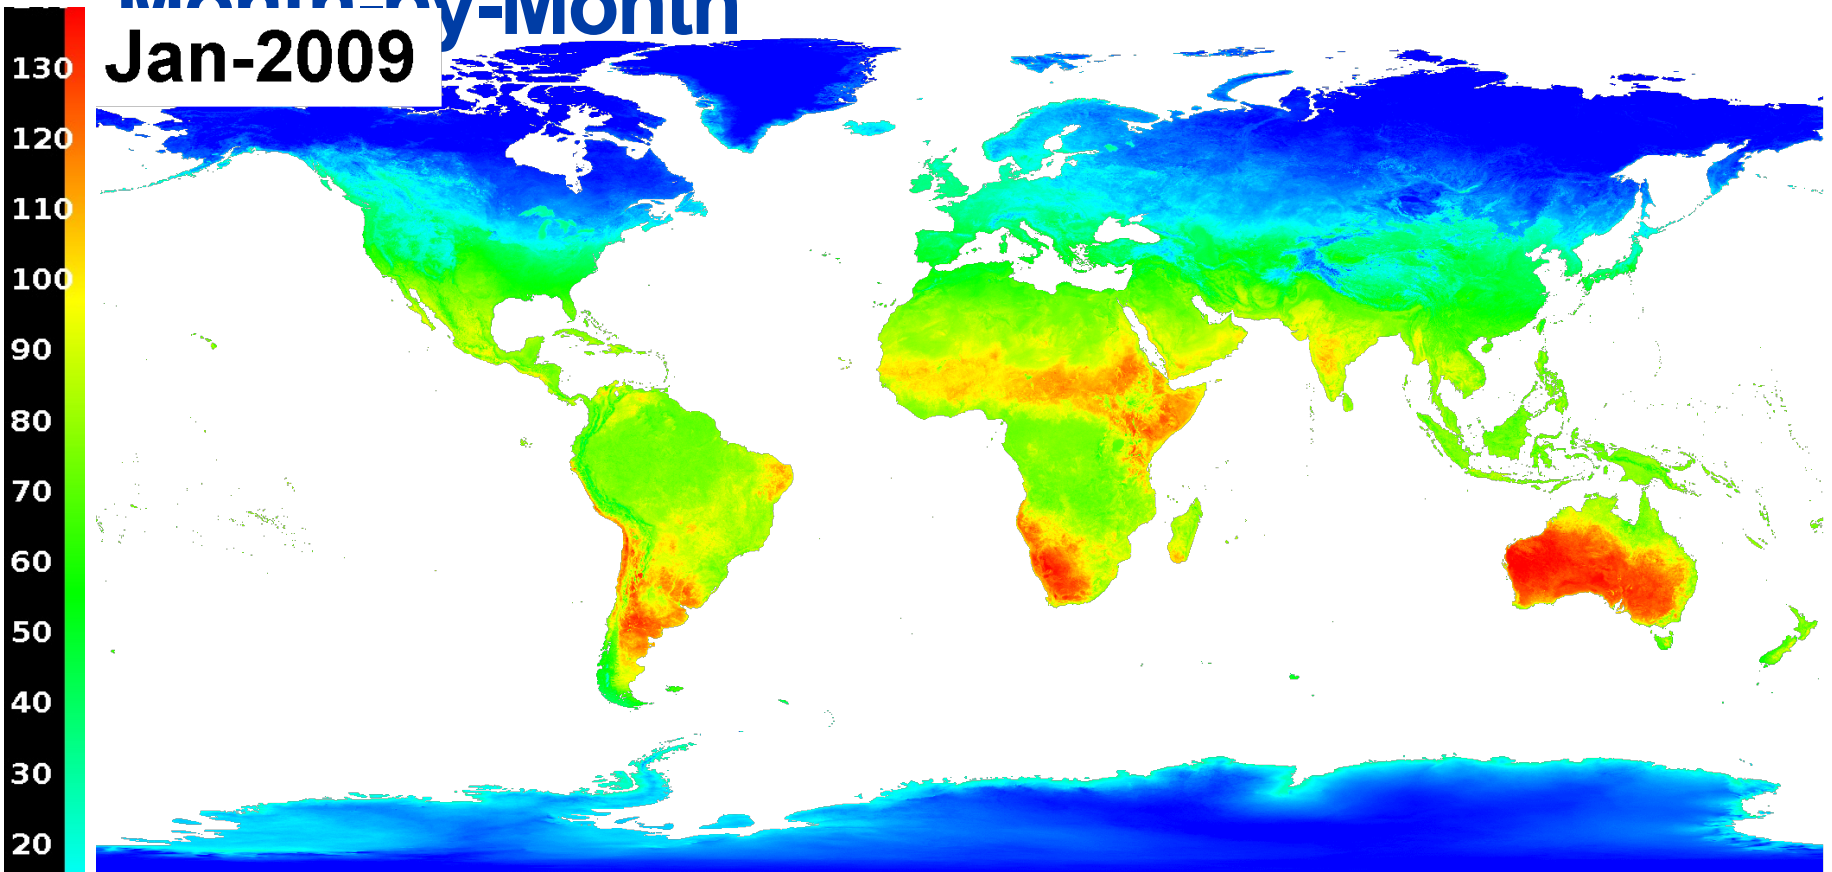


Admin Boundaries

Heat Map From 2009 to 2014

Month-by-Month

Jan-2009

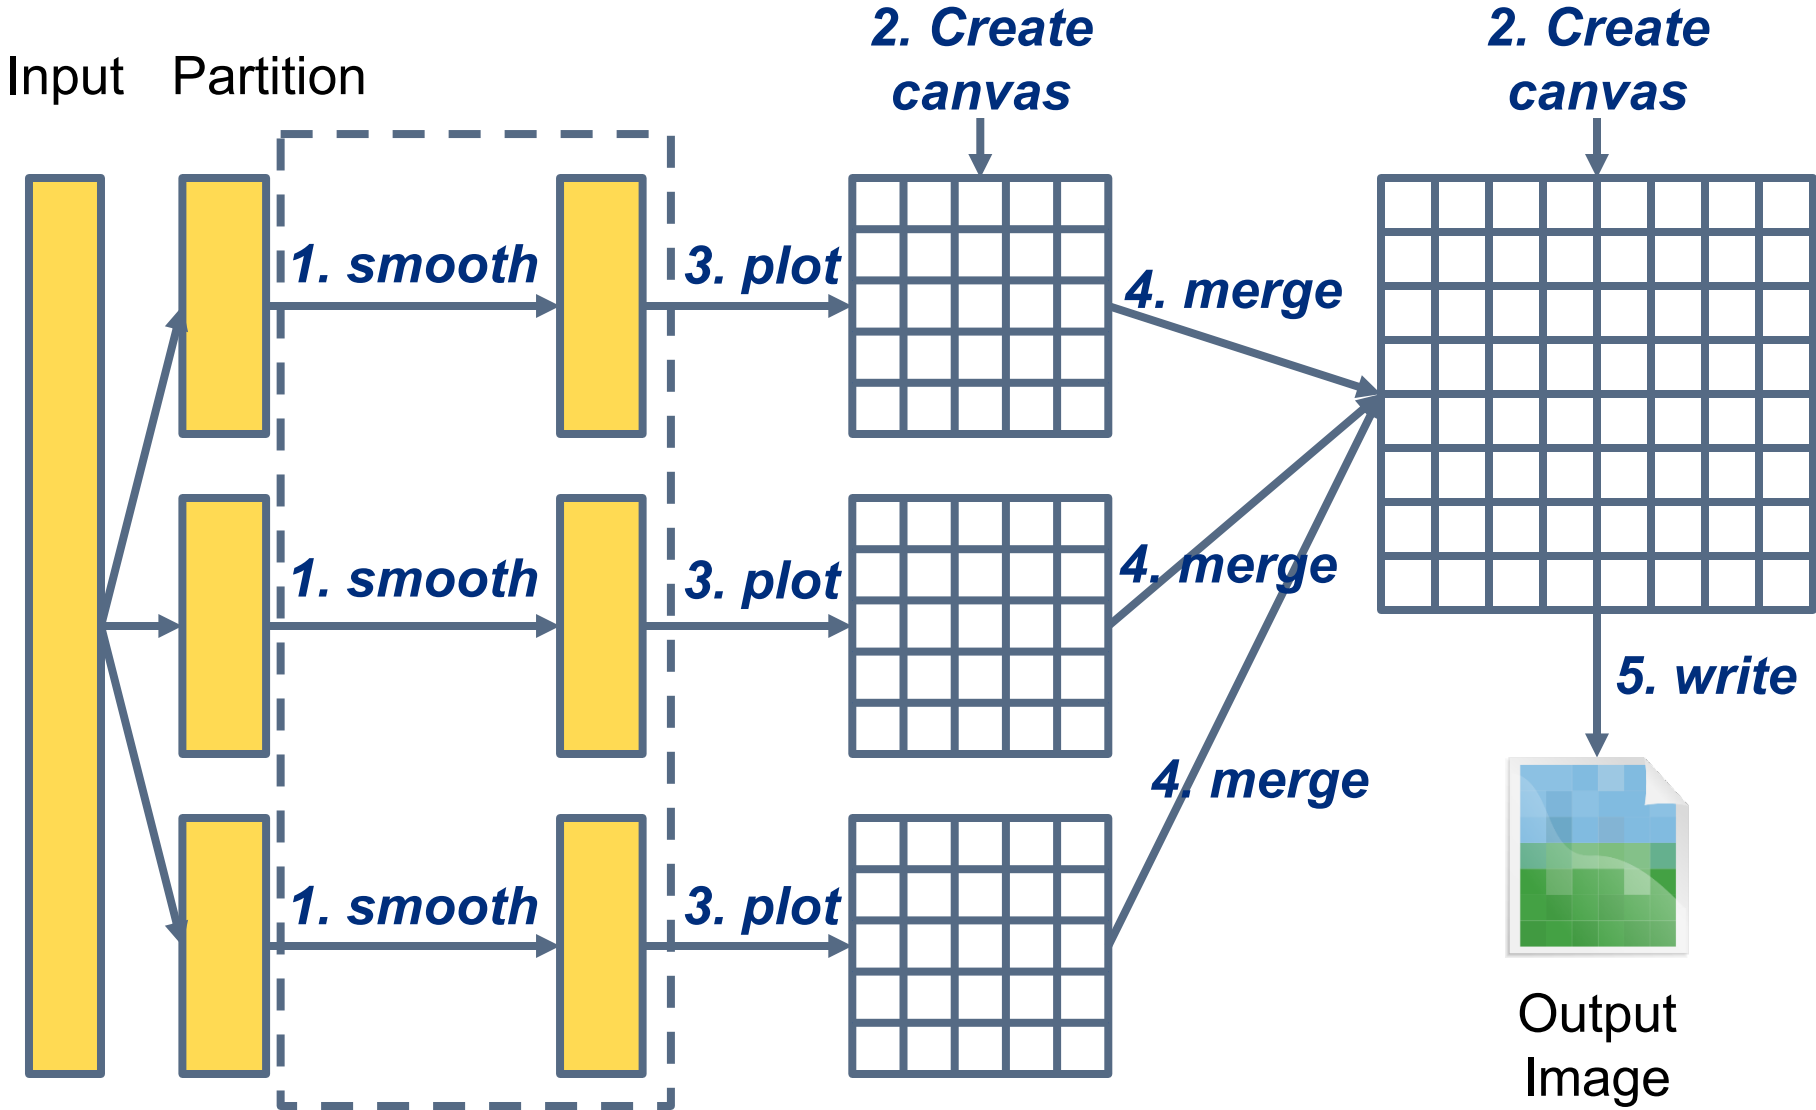


72 Frames × 14 Billion points per frame

Total = **1 Trillion points**


Created in **3 hours** on **10 nodes** instead of **60 hours**

Abstract Visualization



Example: Satellite Data Visualization

1. **Smooth:** Recover holes



2. **Create Canvas:** Initialize a 2D Matrix with zeros

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$


3. **Plot:**
Update the matrix

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 22 & 0 & 0 & 7 & 0 \\ 0 & 0 & 15 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

4. **Merge:**
Matrix addition


$$\begin{bmatrix} & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \end{bmatrix} + \begin{bmatrix} & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \end{bmatrix}$$

5. **Write:**
Generate the image




Example: Road Network Visualization

1. **Smooth:** Merge intersections

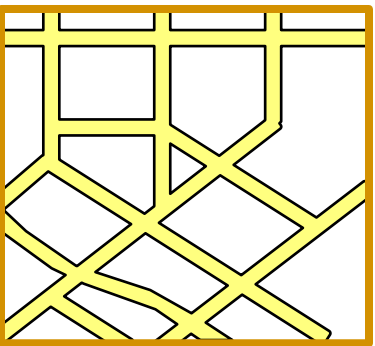


2. **Create Canvas:**


Create a blank image




3. **Plot:** Draw roads as polygons



4. **Merge:** Plot an image on the other



5. **Write:** Encode as PNG and write to file




Single Level Image

Input


Split

caeaeeizester




Split

caeaeeizester




Split

caeaeeizester



Split


caeaeeizester



**Merge
(Overlay)**

Space Partitioning

Input







`create-raster`

`create-raster`


`create-raster`

`create-raster`

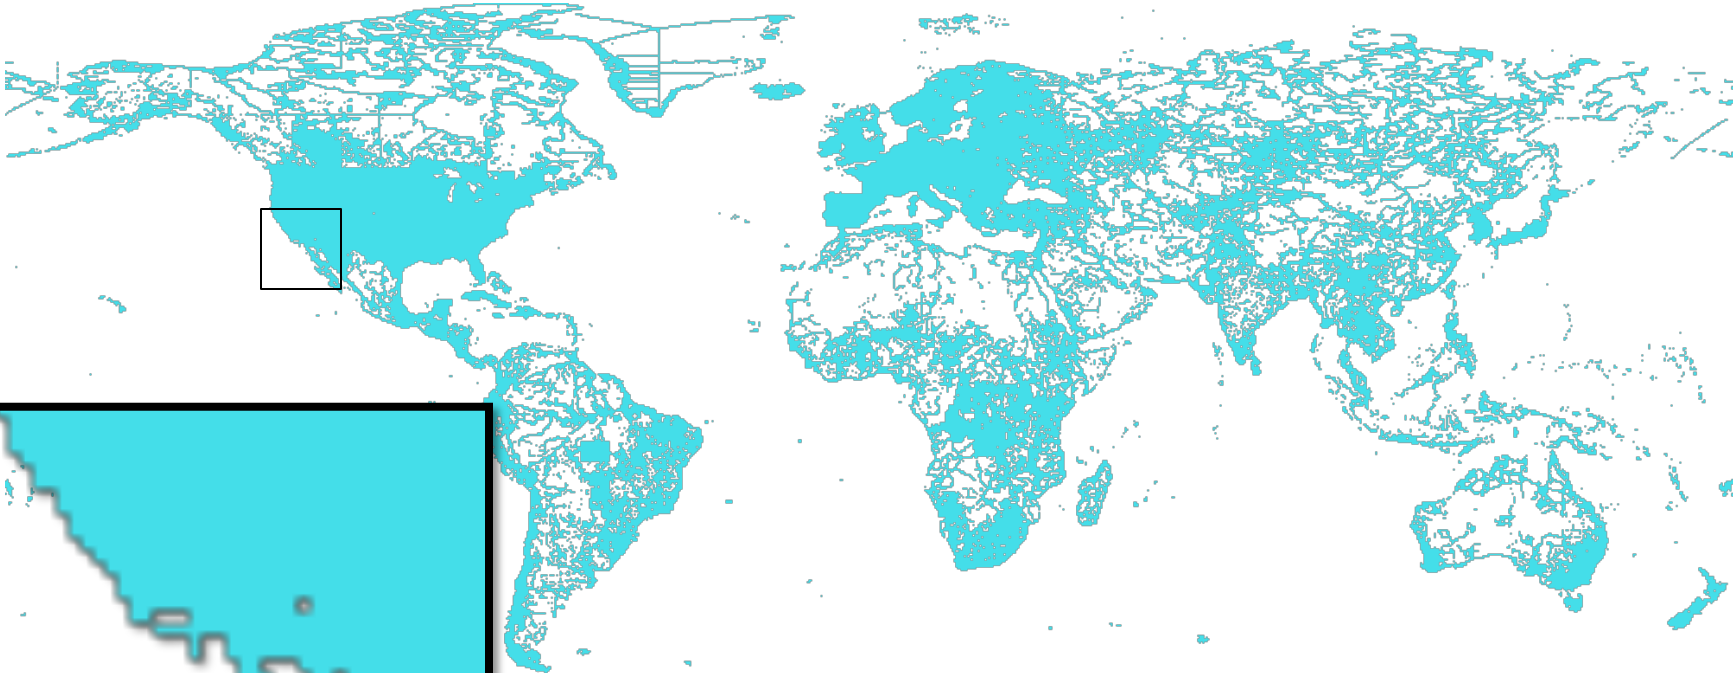


`create-raster`

Merge (Stitch)



Level of Details



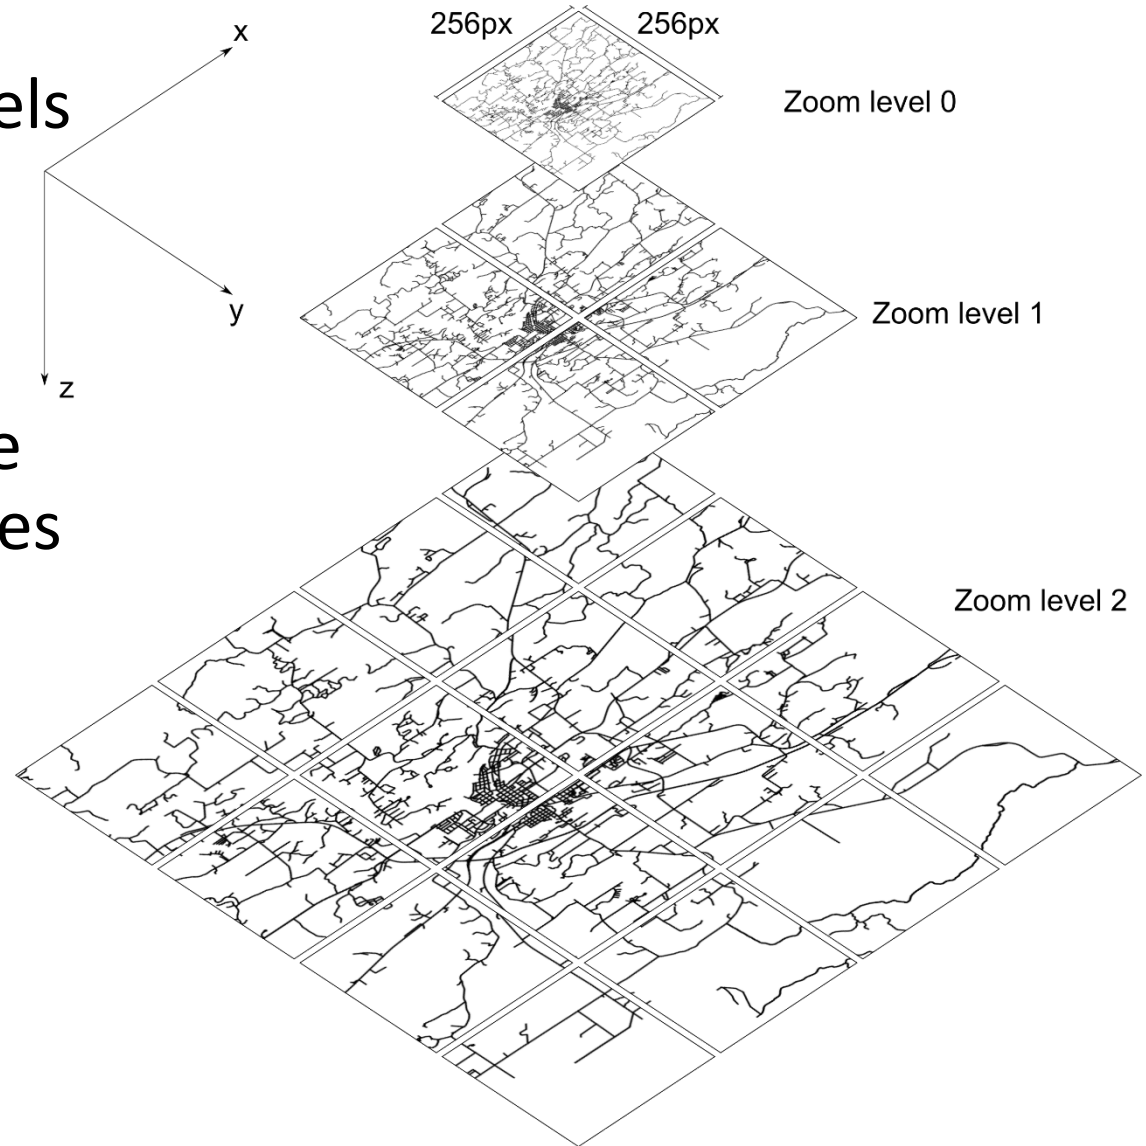


Map of California – 2GB

Generated in **2 minutes** on 10-node cluster instead of **one hour**


Multi-level Image

- Many images at different zoom levels
 - Pan
 - Zoom in/out
 - Fly to
- More details as the zoom level increases
- Number of tiles increases exponentially



Multi-level Visualization

- Abstract multi-level visualization algorithm
- The choice of partitioning technique changes for each zoom level



Beast Architecture



BEAST

Visualization Framework

RDD-based Query Processor

Spatial Partitioner & Load Balancer

In-situ Spark Loaders/Writers

Spatial Data Types

Big
Spatial
Data
Apps

Thank You

Questions?

A Unified Big Data Interface

Unified Big Data Abstraction


Cost Model

Query Optimizer

Query Executor



SparkSQL



MLLib

Spark

Sphinx



Impala



YARN – Resource Manager

HDFS – File System

Language

Applications: SHAHED [ICDE'15] – MNTG [SSTD'13, ICDE'14◇]
TAREEG[SIGMOD'14◇, SIGSPATIAL'14]



Language
Pigeon [ICDE'14◇]



Visualization
[VLDB'15◇, ICDE'16]

Operations

Basic operations – CG_Hadoop
[SIGSPATIAL'13, TSAS★]

MapReduce

Spatial File Splitter
Spatial Record Reader

Indexing

Grid – R-tree – R+-tree – Quad tree
[VLDB'15]

ST-Hadoop [TODS★]

★ Under review

◇ Demo paper

Language (Pigeon)

- Hides the complexity of the system with a high level language
- OGC standard used by Oracle Spatial and PostGIS
- Extends Pig Latin with OGC-compliant primitives
 - Spatial data types (e.g., Polygon)
 - Basic operations (e.g., Area)
 - Spatial predicates (e.g., Touches)
 - Spatial analysis (e.g., Union)
 - Spatial aggregate functions (e.g., Convex Hull)



Spatial Data Types

Data Loading

```
lakes = LOAD 'lakes' AS (id:int, area:polygon);
```

Range Query

```
houses_in_range =  
  Filter houses BY  
  Overlap(house_loc, range);
```

KNN


```
nearest_houses =  
  KNN houses WITH_K=100  
  USING DistanceTo(house_loc,  
  query_loc);
```

Spatial Join

```
lakes_states = Join lakes BY lakes_boundary  
  states BY states_boundary Predicate = Overlap
```

Spatio-temporal Indexing

Applications: SHAHED [ICDE'15] – MNTG [SSTD'13, ICDE'14◇]
TAREEG[SIGMOD'14◇, SIGSPATIAL'14]



Spatial Hadoop

VLDB'13◇

ICDE'15

Language

Pigeon [ICDE'14◇]



Visualization

HadoopViz[VLDB'15◇]

Operations

Basic operations – CG_Hadoop
[SIGSPATIAL'13, TSAS★]

MapReduce

Spatial File Splitter
Spatial Record Reader

Indexing

Grid – R-tree – R+-tree – Quad tree
[VLDB'15]

ST-Hadoop [TODS★]

A. Eldawy, L. Alarabi, M. F. Mokbel. "ST-Hadoop: A MapReduce Framework for Spatial and Spatio-temporal Data" Submitted to **ACM TODS**

Multiresolution Spatio-temporal Index

2012

2013

Yearly Indexes

jan

feb

...

dec

jan

feb

...

dec

jan

Monthly Indexes

1

2

...

366

1

2

...

365

1

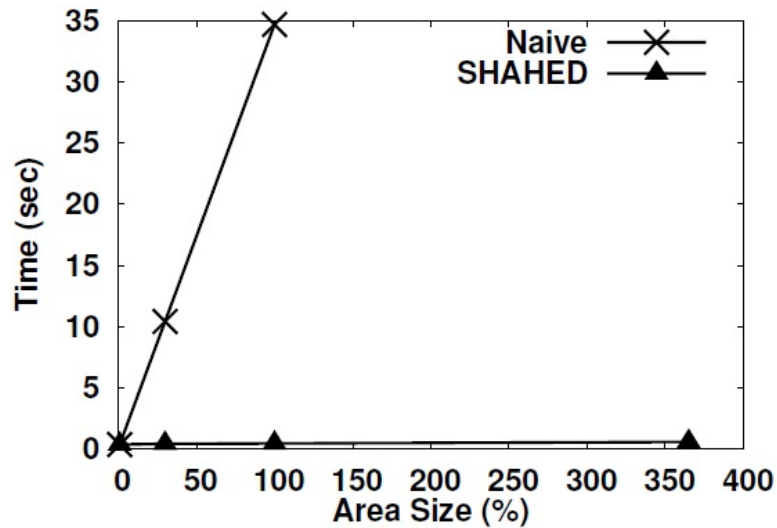
2

...

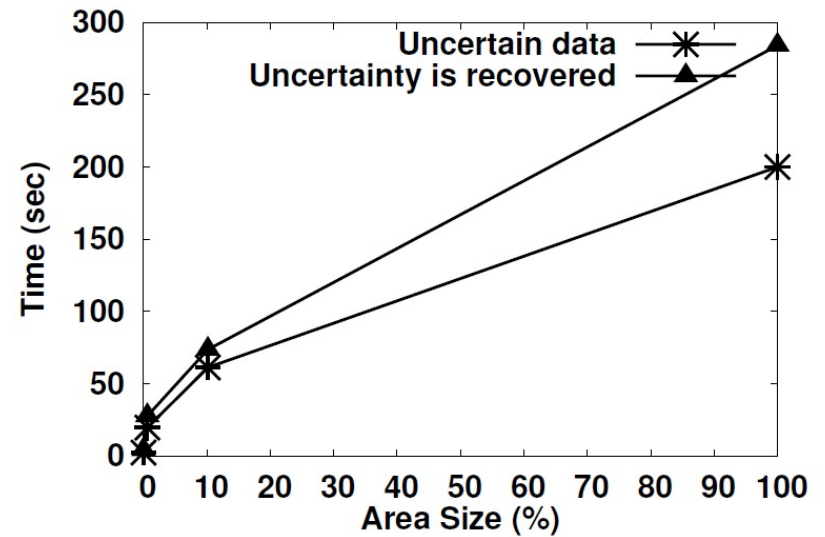
31

Daily Indexes

Performance of SHAHED




(c) Selection Query

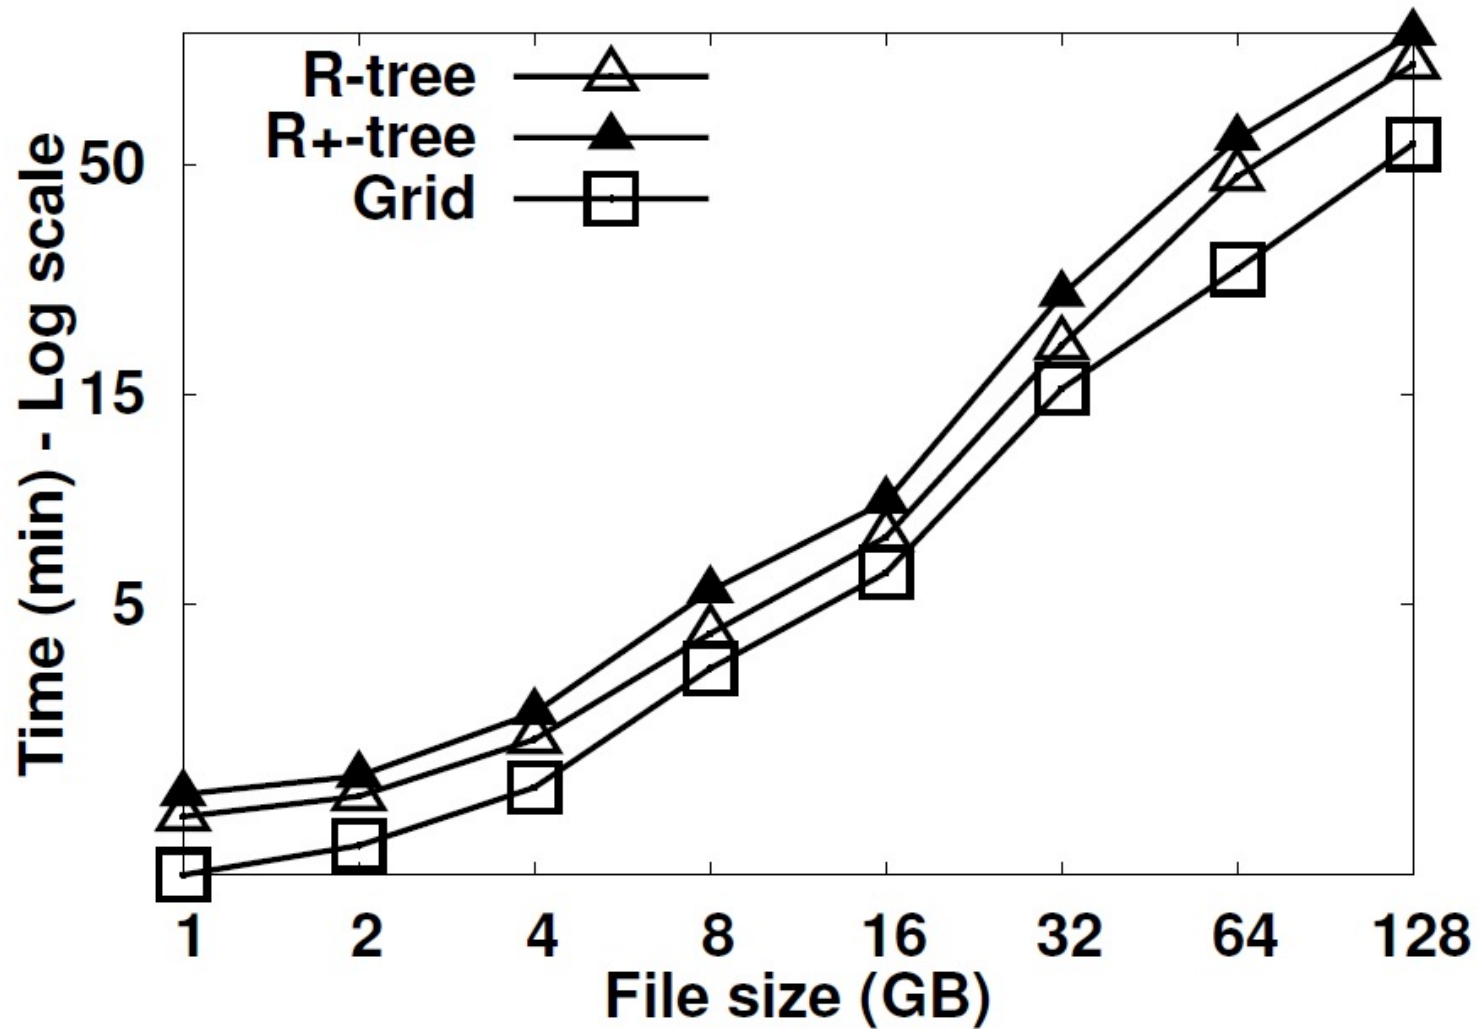


(d) Visualization

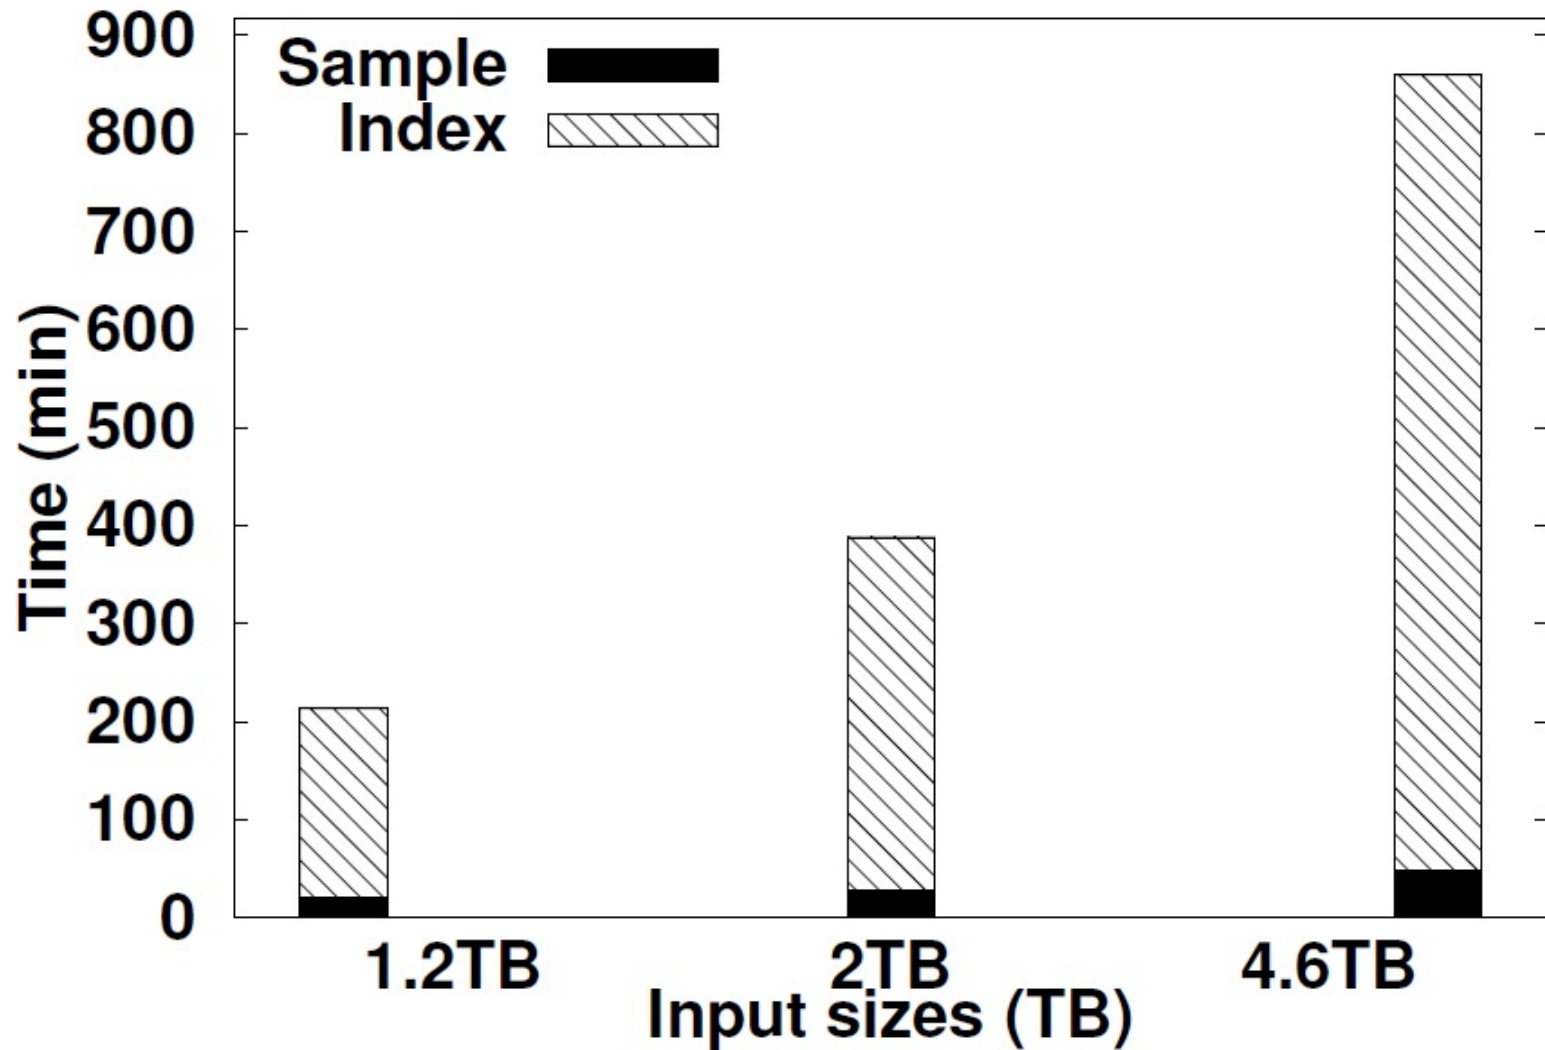
Reference Point



Index building



Index Building for NASA Data




Related Work

- Most techniques for spatial data processing in Hadoop use Hadoop as a blackbox
 - RQ, KNN and SJMR [Zhang et al'09]
 - R-tree construction [Cary et al'09]
 - KNN Join [Lu et al'12, Zhang et al'12]
 - RNN [Akdogan et al'10]
 - ANN [Wang et al'10]
- MD-HBase [Nishimura et al'11]
 - Framework for multi-dimensional data processing
 - Based on HBase, a key-value store on HDFS
 - Does not support MapReduce programming

N Map plan – SpatialHadoop


Indexed Input
File(s)



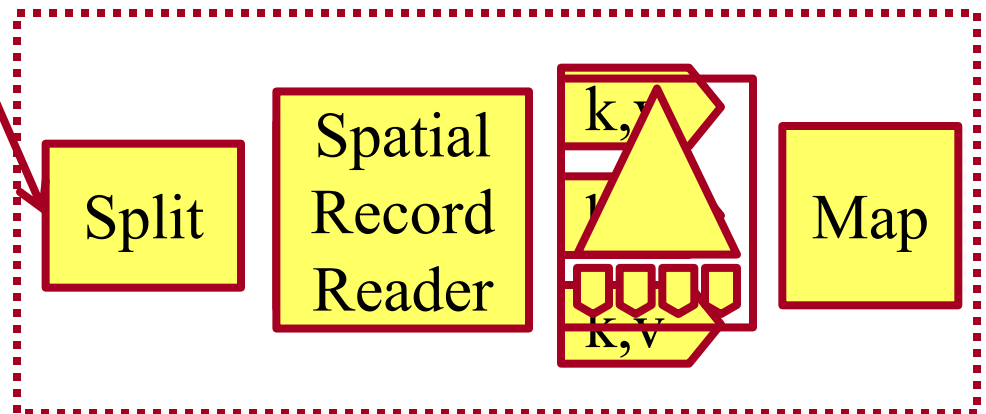
Number
of splits



Map task



⋮
Map task



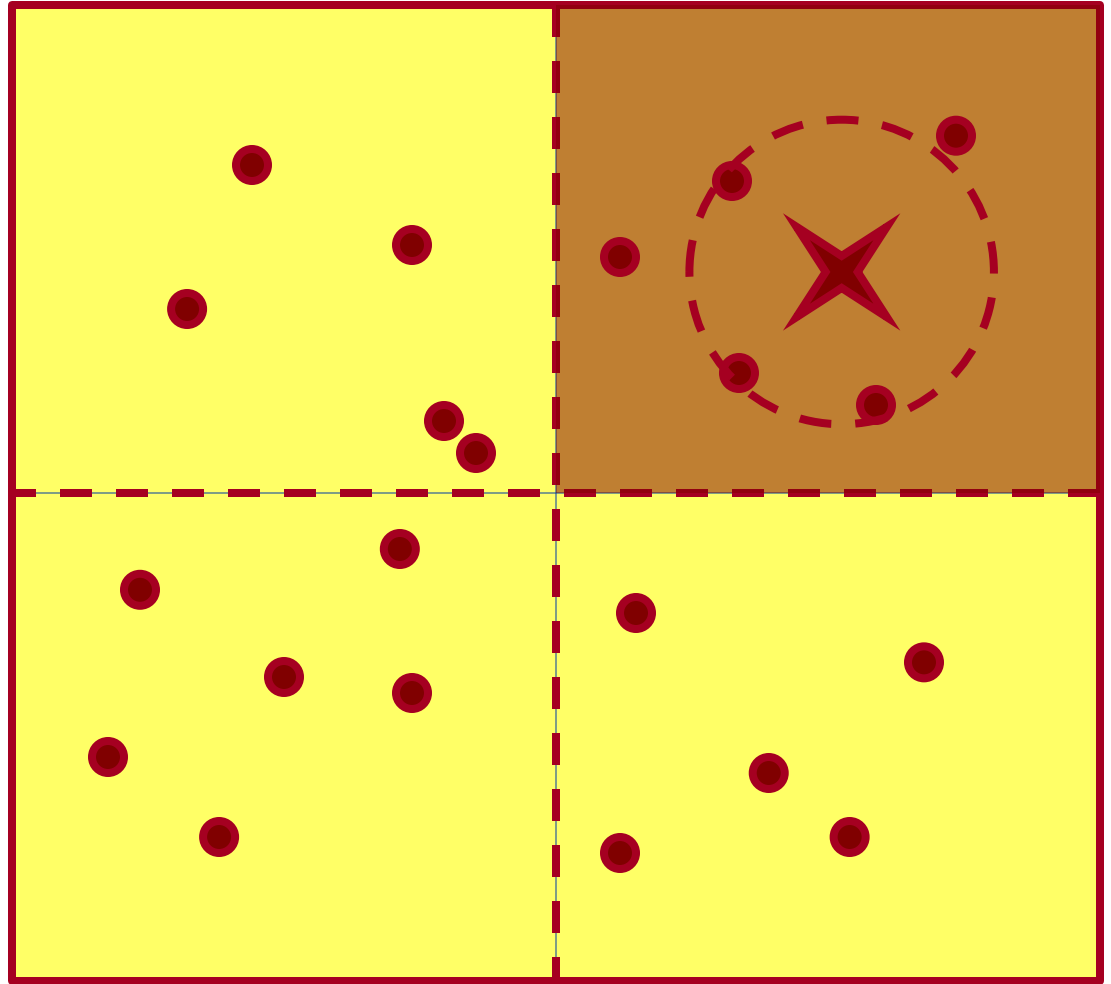
KNN

SpatialFileSplitter
selects the block that
contains the query
point

Map function performs
kNN in the selected
block

Answer is tested for
correctness

✓ Answer is correct




k=3

KNN

First iteration runs as before and result is tested for correctness

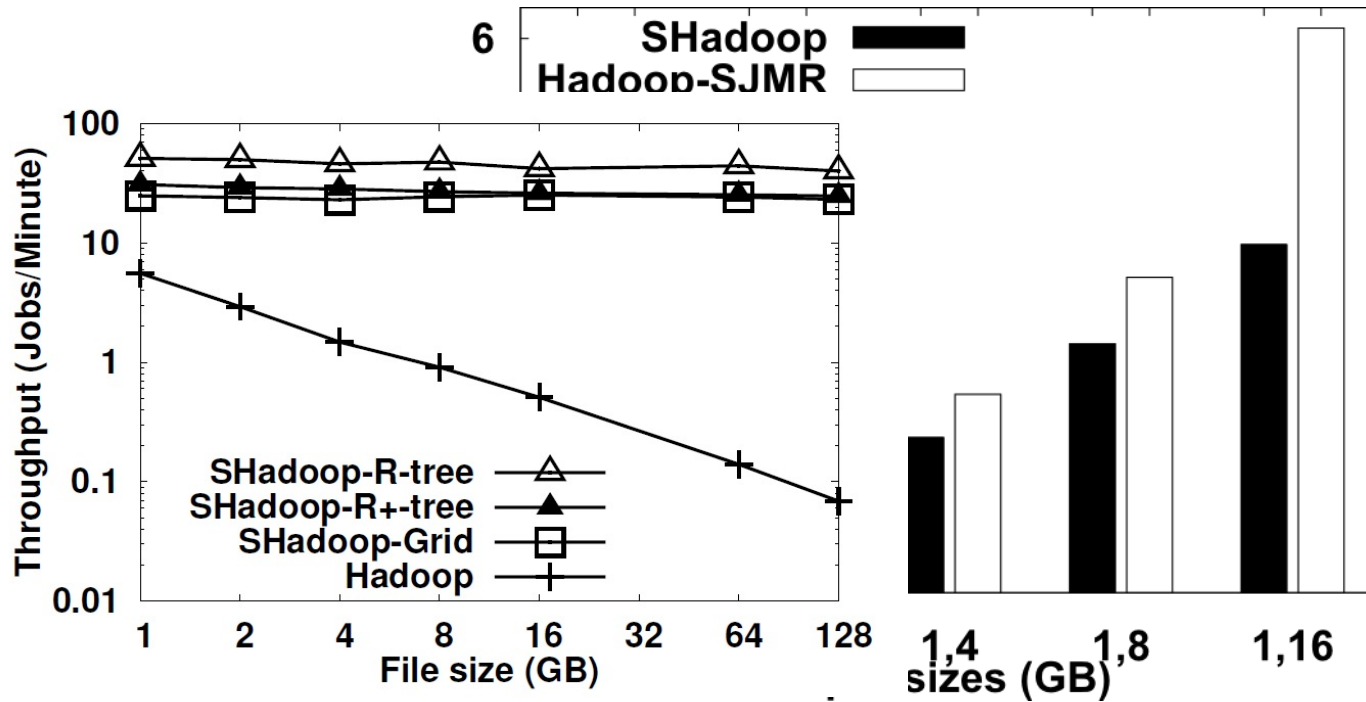
✗ Answer is incorrect

Second iteration processes other blocks that might contain an answer

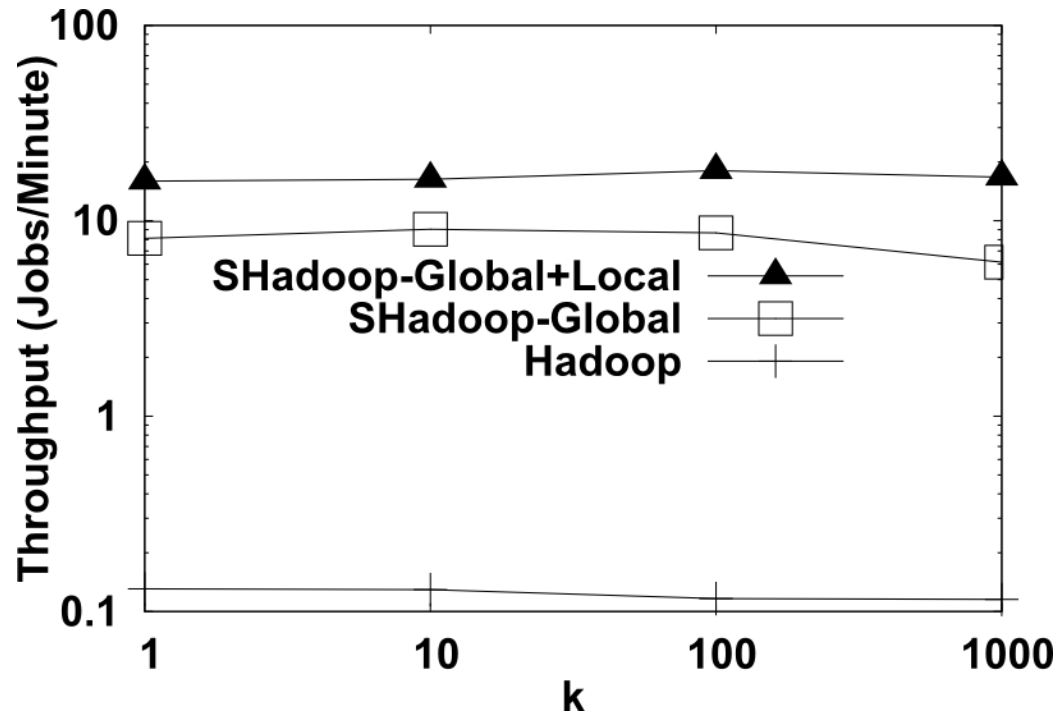


$k=3$

Range query



K Nearest Neighbor



Preliminary Results

