

# GeoTrellis

## Adding Geospatial Capabilities to Spark

Rob Emanuele  
lossyrob@gmail.com



@lossyrob



LocationTech

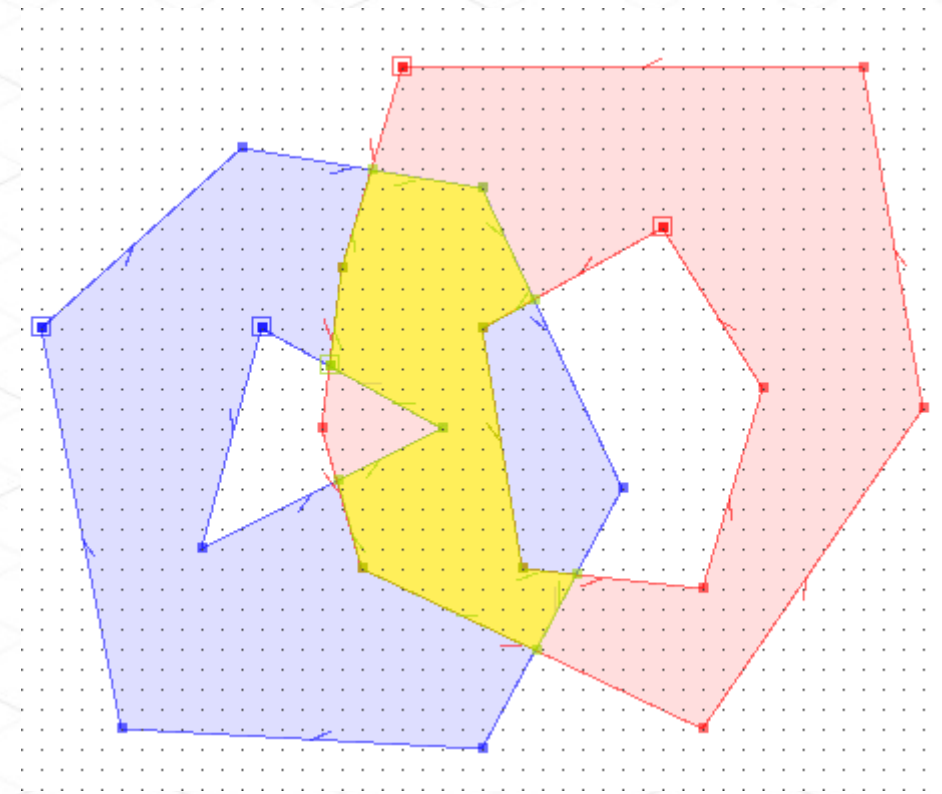
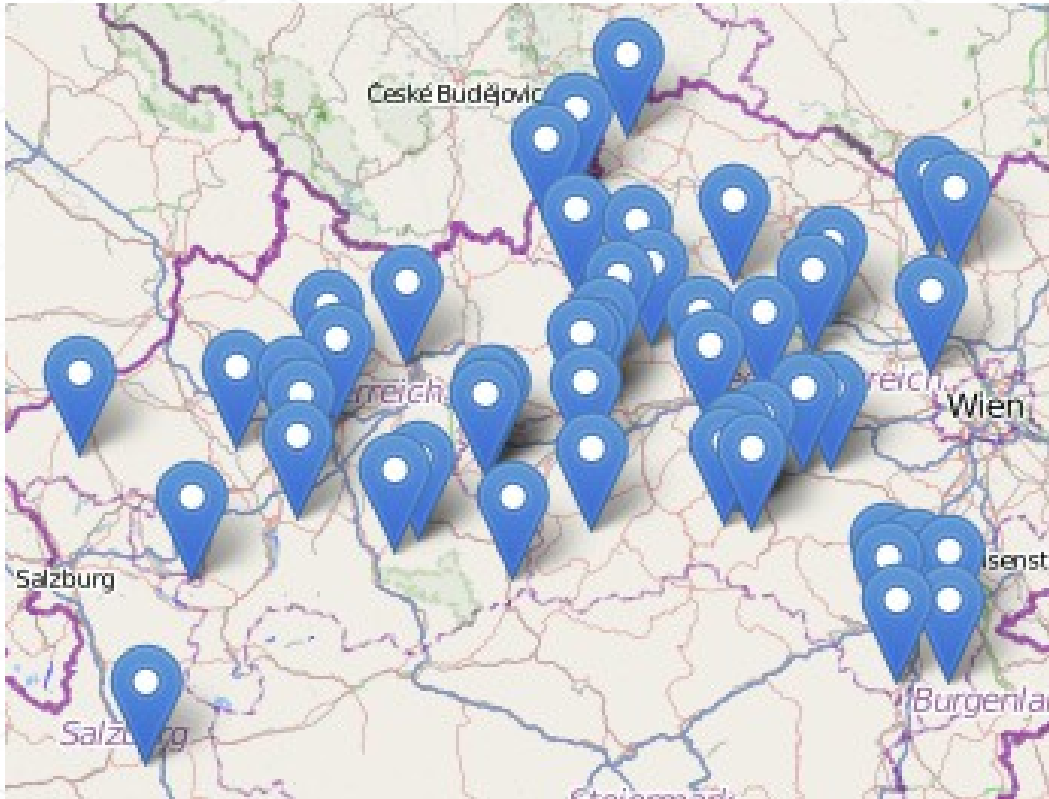


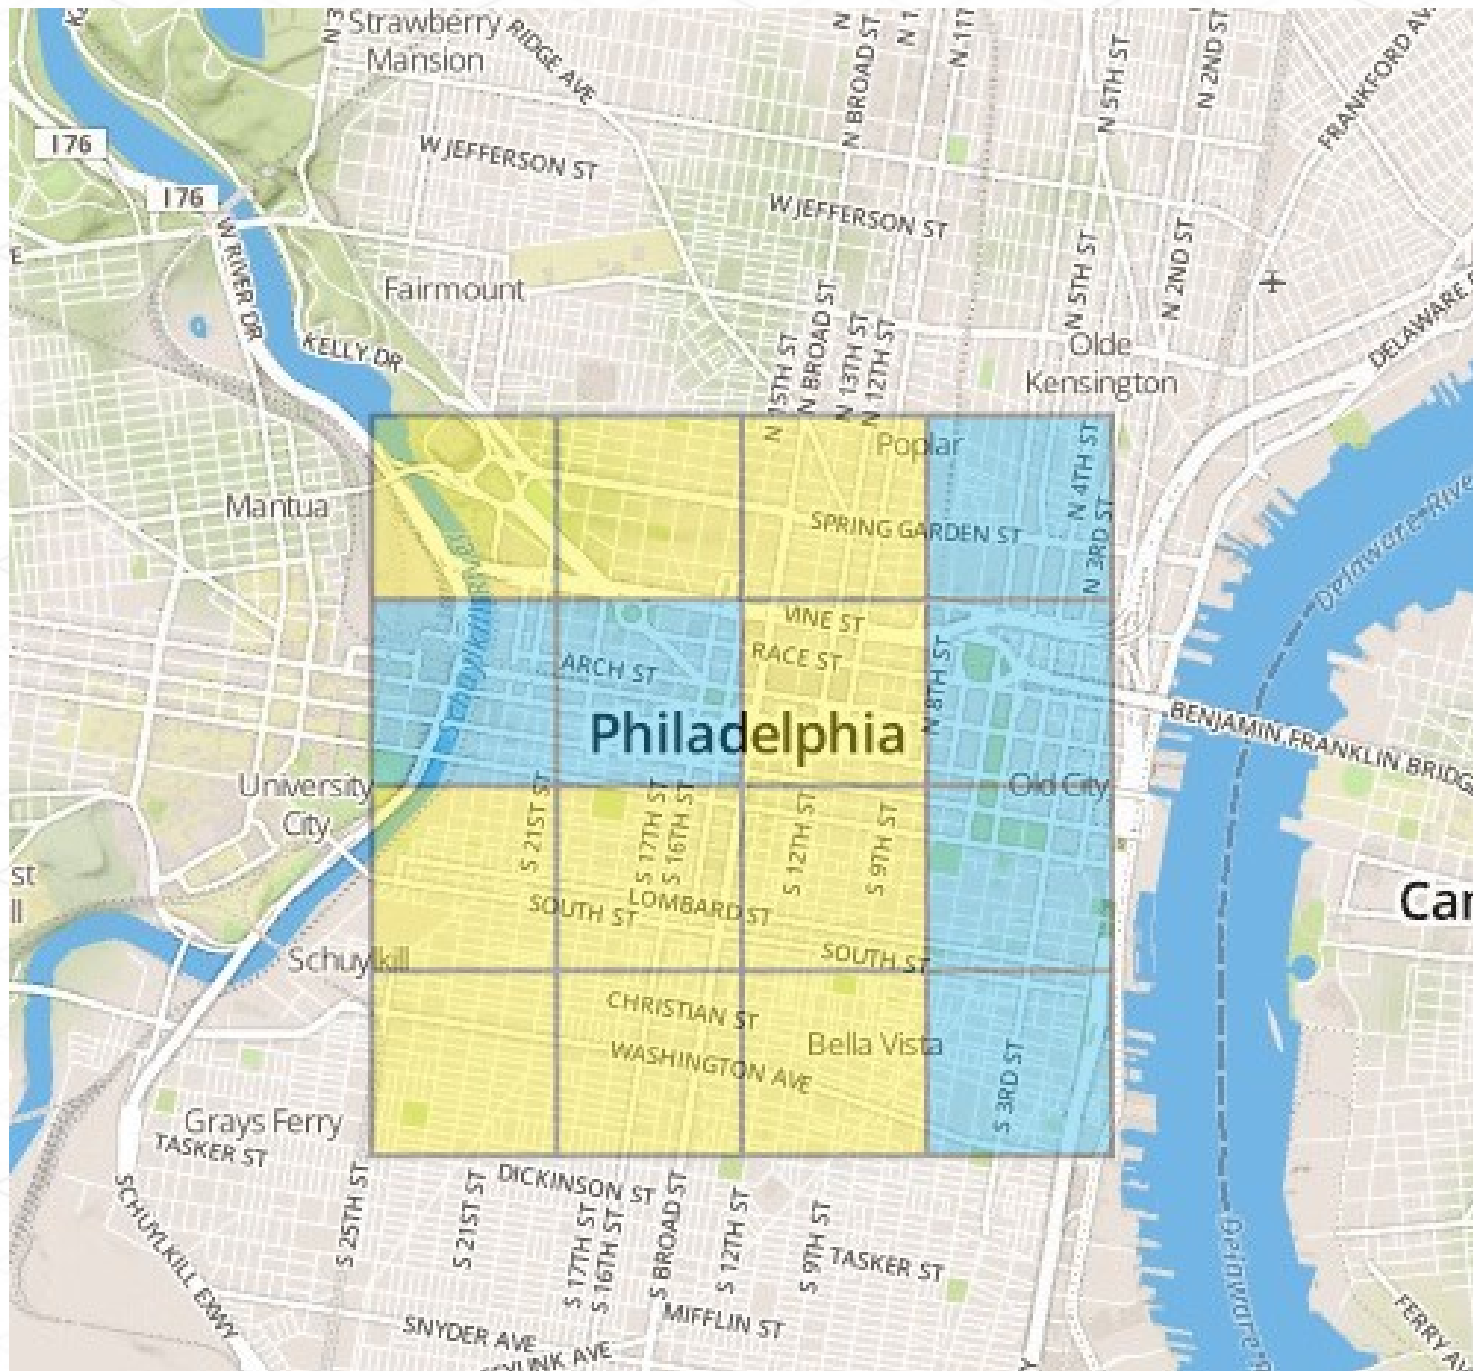
# GeoTrellis

## What is GeoTrellis?

- Geospatial data types\IO (raster & vector)
- Raster operations (Map Algebra)
- Utilities for creating web services

# Vector Data





# NED (National Elevation Dataset)







# GeoTrellis

## What is GeoTrellis?

- Geospatial data types\IO (raster & vector)
- Raster operations (Map Algebra)
- Utilities for creating web services



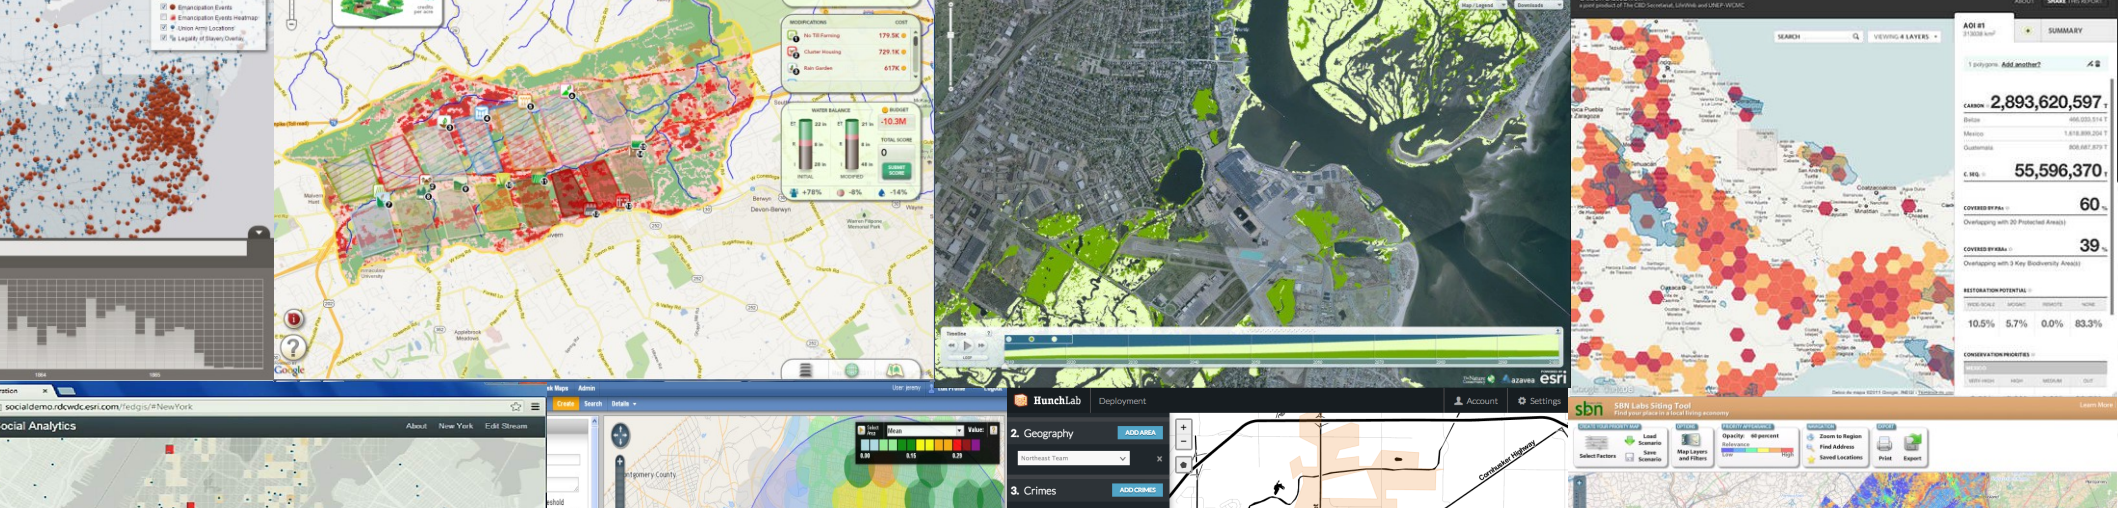
# GeoTrellis

## What is GeoTrellis?

- Geospatial data types\IO (raster & vector)
- Raster operations (Map Algebra)
- Utilities for creating web services

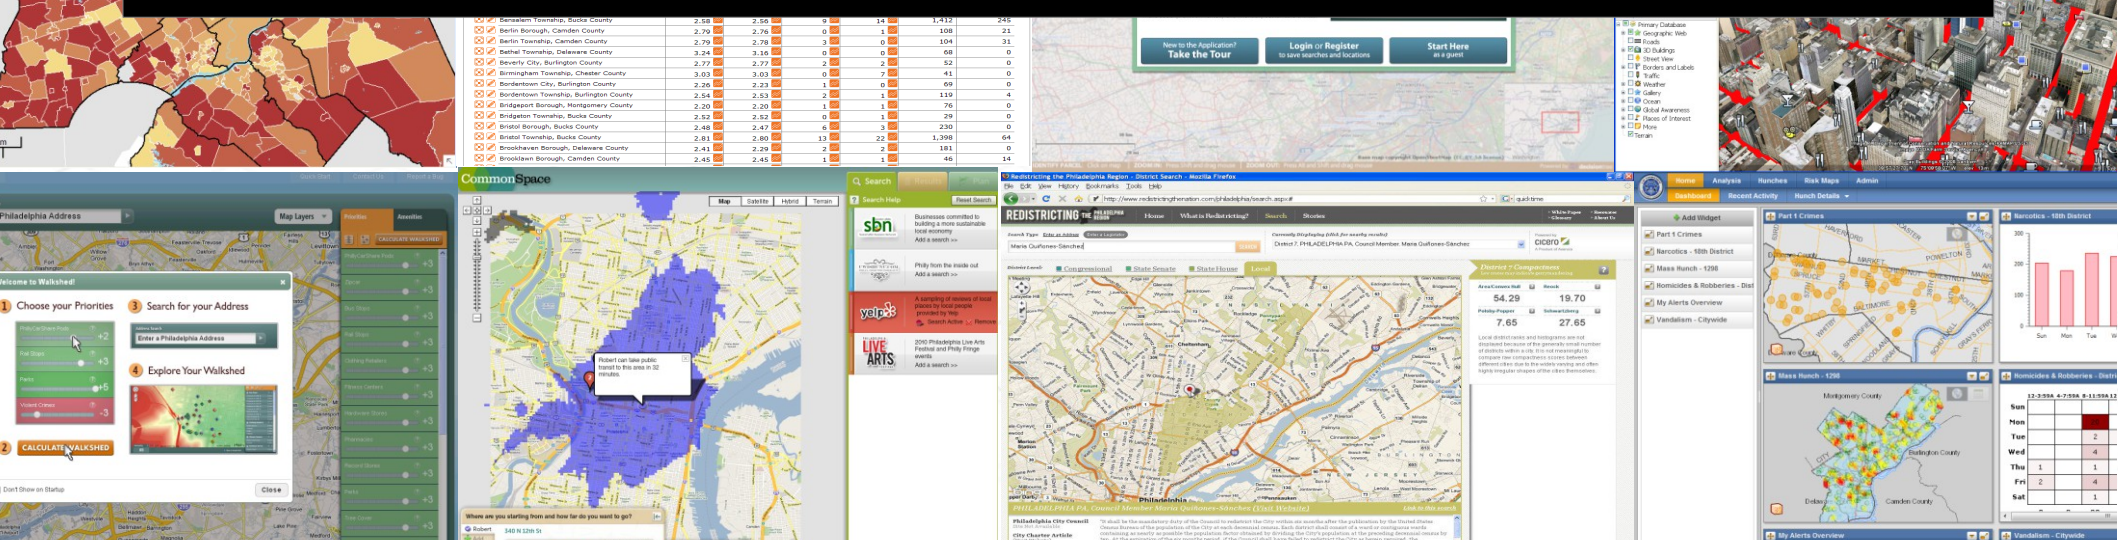
**geospatial** with **scala**






# azavea

beyond dots on a map





**sbn** SBN Labs Siting Tool  
Find your place in a local living economy

Learn More | Login | Tour

**CREATE YOUR PRIORITY MAP**  
Select Factors | Load Scenario | Save Scenario

**OPTIONS**  
Map Layers and Filters

**PRIORITY APPEARANCE**  
Opacity: 60 percent  
Relevance: Low to High

**NAVIGATION**  
Zoom to Region | Find Address | Saved Locations

**EXPORT**  
Print | Export


IDENTIFY PARCEL: Click on map | ZOOM IN: Press Shift and drag mouse | ZOOM OUT: Press Alt and Shift and drag mouse

Base map copyright OpenStreetMap (CC-BY-SA license)

Powered by **decisiontree**


# Business Siting





6183 x 4992

118 MB



4598 x 4867

86 MB



1770271 x 910139

5.8 TB





# DigitalGlobe™





GeoTrellis

+



## GeoTrellis v0.9:

- ▣ Parallel operations across tiles
- ▣ Parallel execution of operations
- ▣ Basic cluster capabilities with Akka clustering



GeoTrellis

+



## What's missing?

- ▣ Sharding raster data across the cluster
- ▣ Caching operation results across cluster
- ▣ HDFS support
- ▣ Advanced Fault tolerance
- ▣ Advanced Scheduling
- ▣ ...






+




Level 1




Level 2



Level 3



Level 2



Partitioned across two map files, one for each row

Key	Value
0	Tile 0
1	Tile 1
2	Tile 2
3	Tile 3

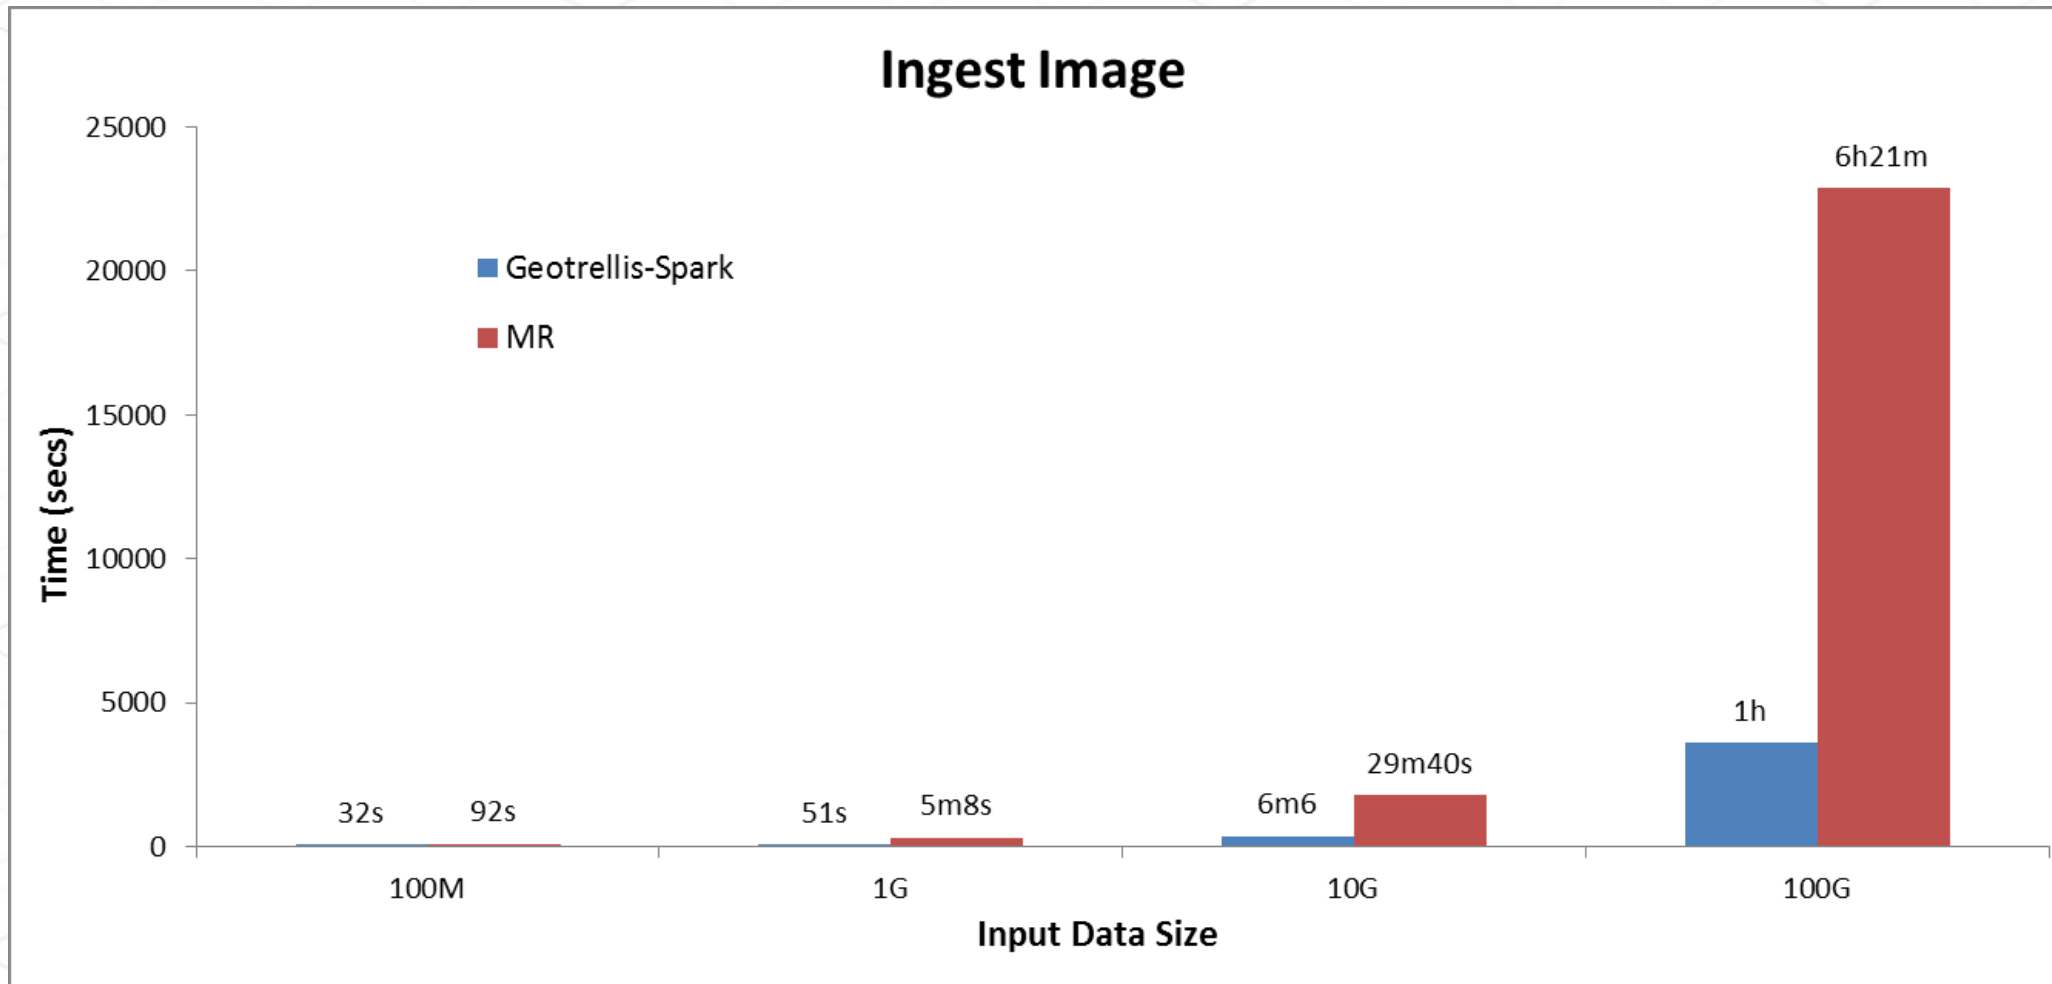
Key	Value
4	Tile 4
5	Tile 5
6	Tile 6
7	Tile 7

**Data Types**  
 Key: Long  
 Value: Raster


# RasterRDD

- Custom Partitioner based on TileID so that each worker works with one MapFile in HDFS.
- Defines raster operations in terms of Spark transformation and actions like mapPartitions, zipPartition, etc.
- Operations use logic already written and tested in core GeoTrellis, but use Spark for distributing those operations over tiles.

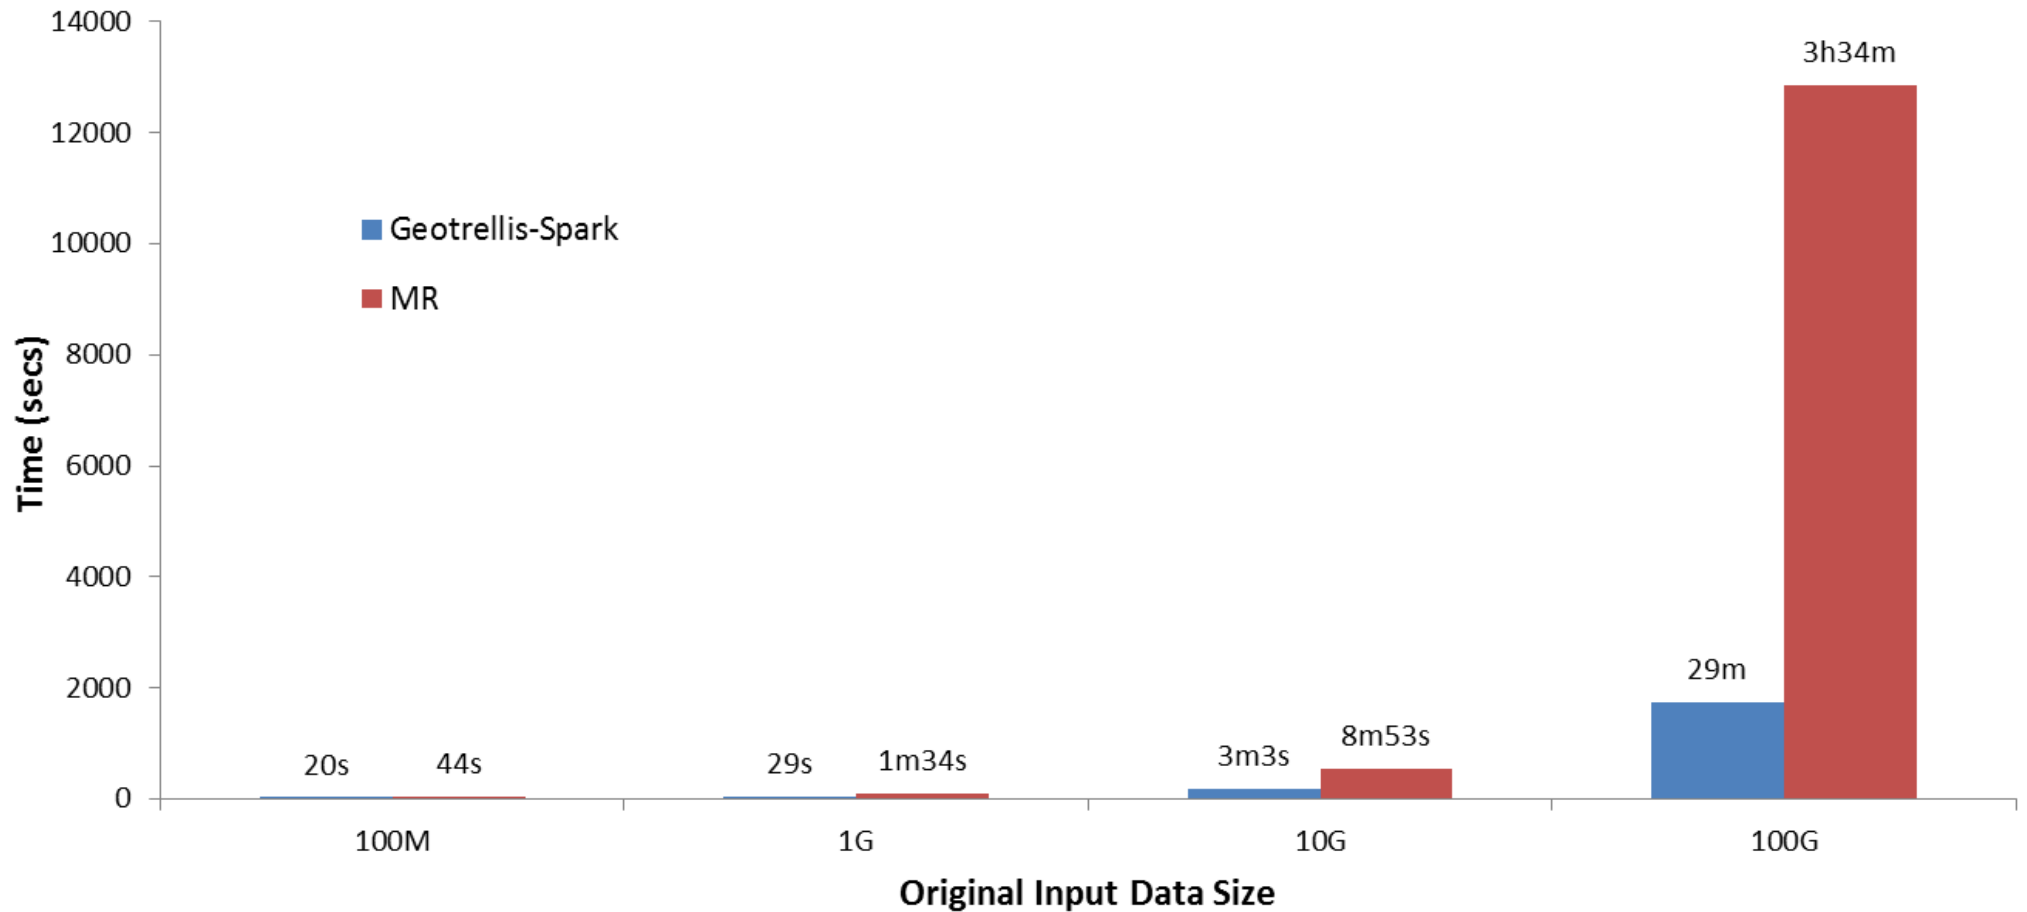




## Build Pyramid

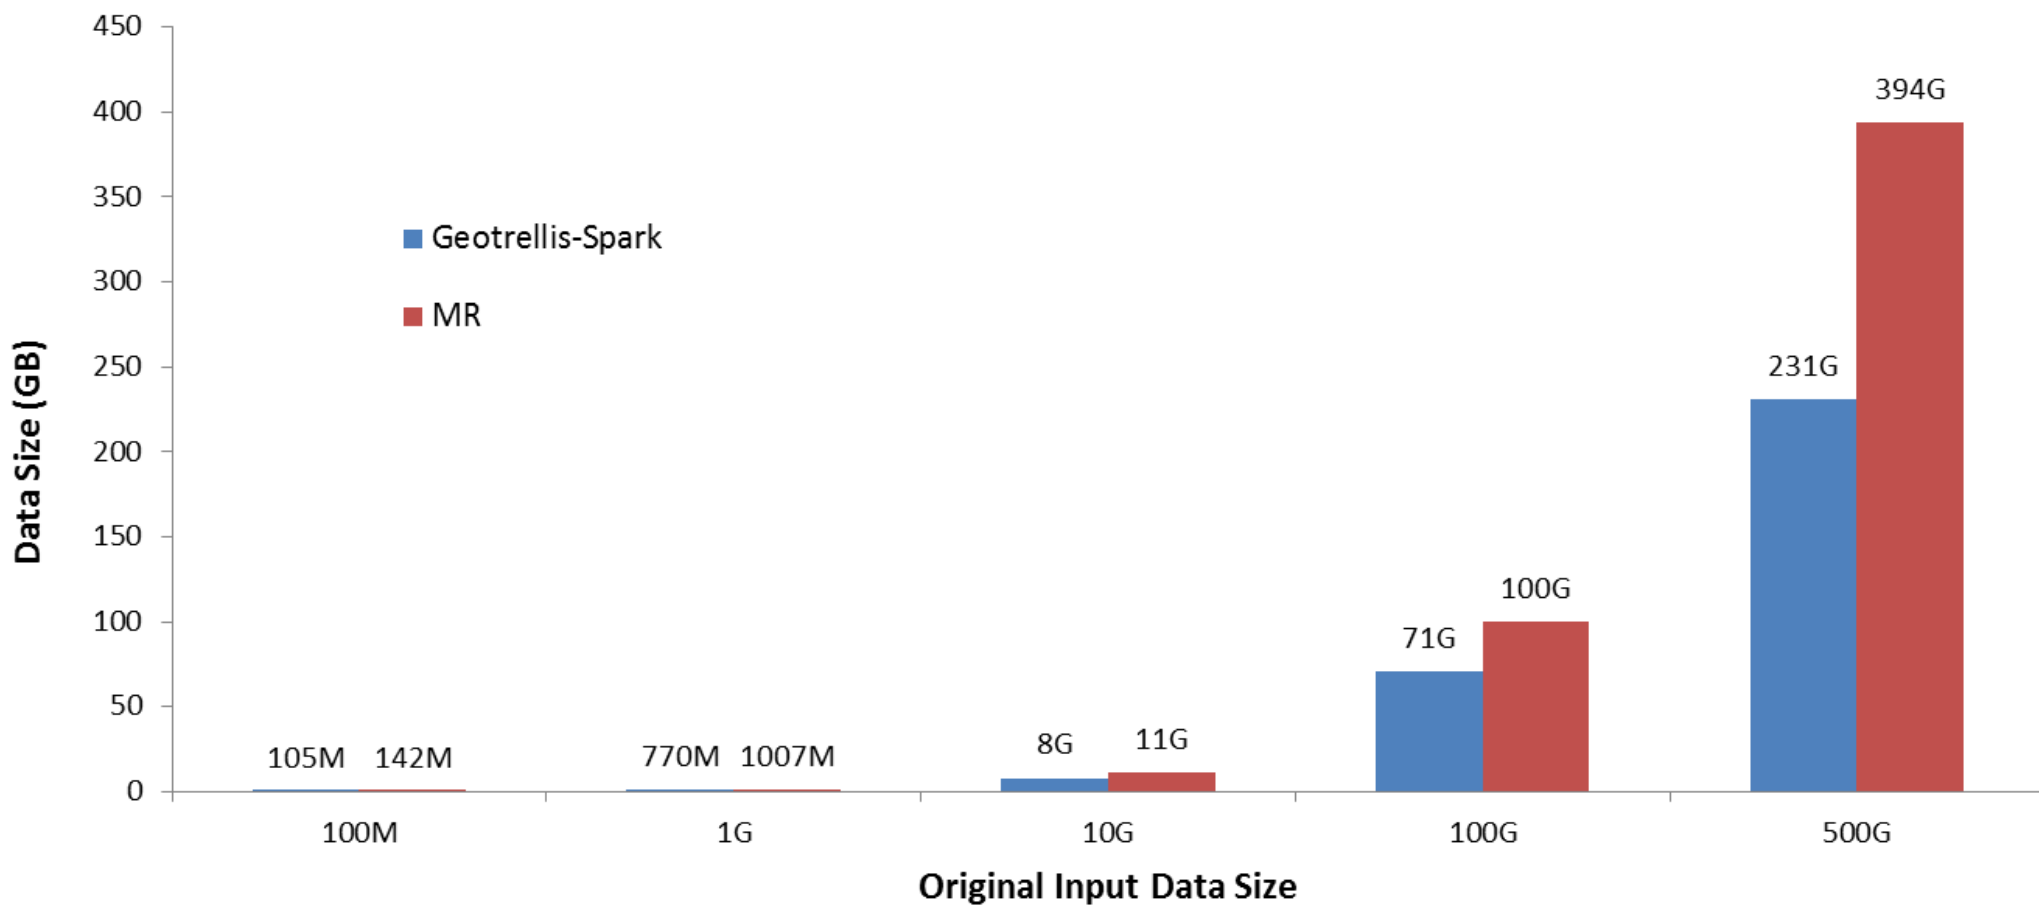


## Add Rasters

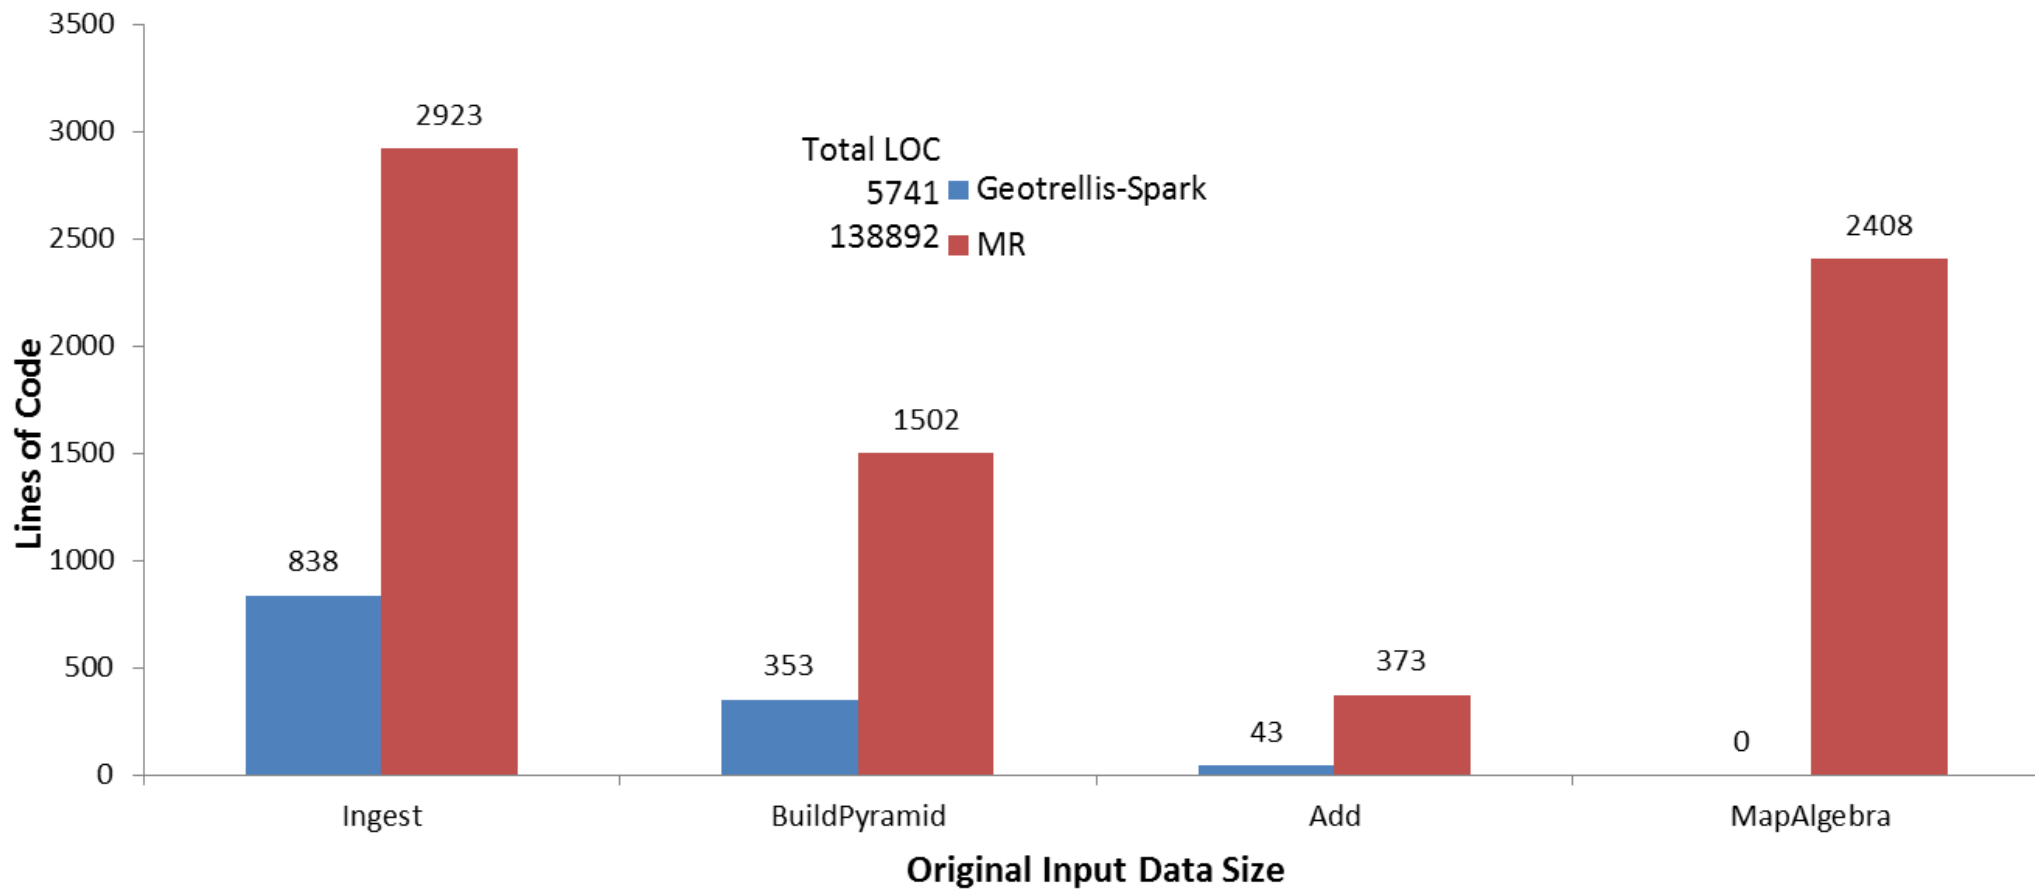




## Disk Space



## Code Size



# What's next?

- More operations: Focal operations, zonal operations
- Creating tile serving mechanisms for web services.
- Vector operations on spark?

**Spark**  
Lightning-Fast Cluster Computing

+



GeoTrellis

=

geospatial at **scale**





# GeoTrellis



geotrellis.io  
[github.com/geotrellis/geotrellis](https://github.com/geotrellis/geotrellis)

**IRC: #geotrellis on freenode**