University of Washington Astronomy
Survey Science Group

Astronomical Image Processing with Hadoop

Keith Wiley, Andrew Connolly, Simon Krughoff, Jeff Gardner,
Magdalena Balazinska, Bill Howe, YongChul Kwon, Yingyi Bu

NSF Cluster Exploratory (CluE) grant IIS-0844580
NASA grant 08-AISR08-0081

Future astronomical surveys will generate 10s of TBs of image data and detect millions of sources per night.

- **Example:** LSST* (2015-2025)
- 8.4m mirror
- 3.2 Gpixel camera
- Half sky every three nights
- 30 TBs per night
- 60 PBs total
- 1000s of exposures per location

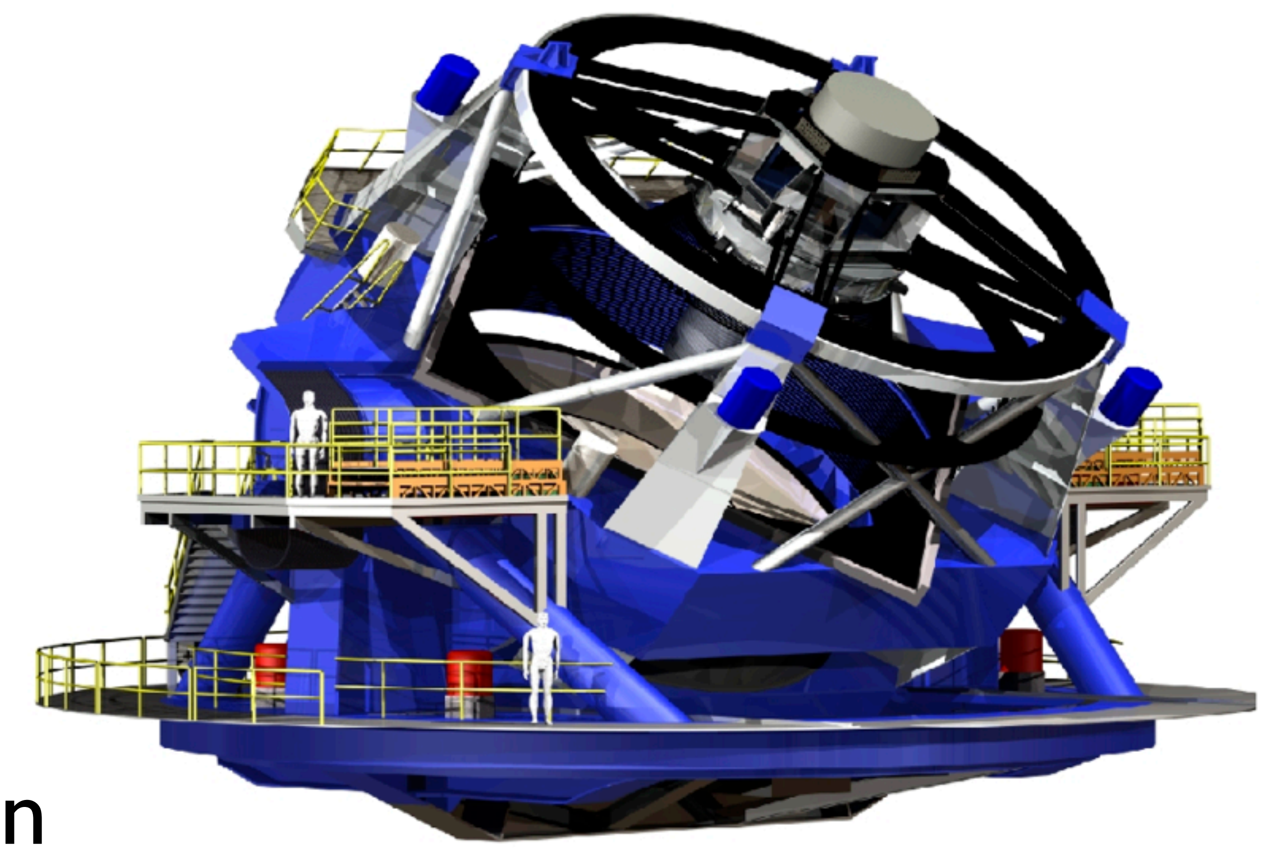

Astronomers will need to analyze and detect moving/transient sources in real time.

This challenge is beyond desktop capabilities.

* Large Synoptic Survey Telescope

Massively parallel databases and computing clusters are required.

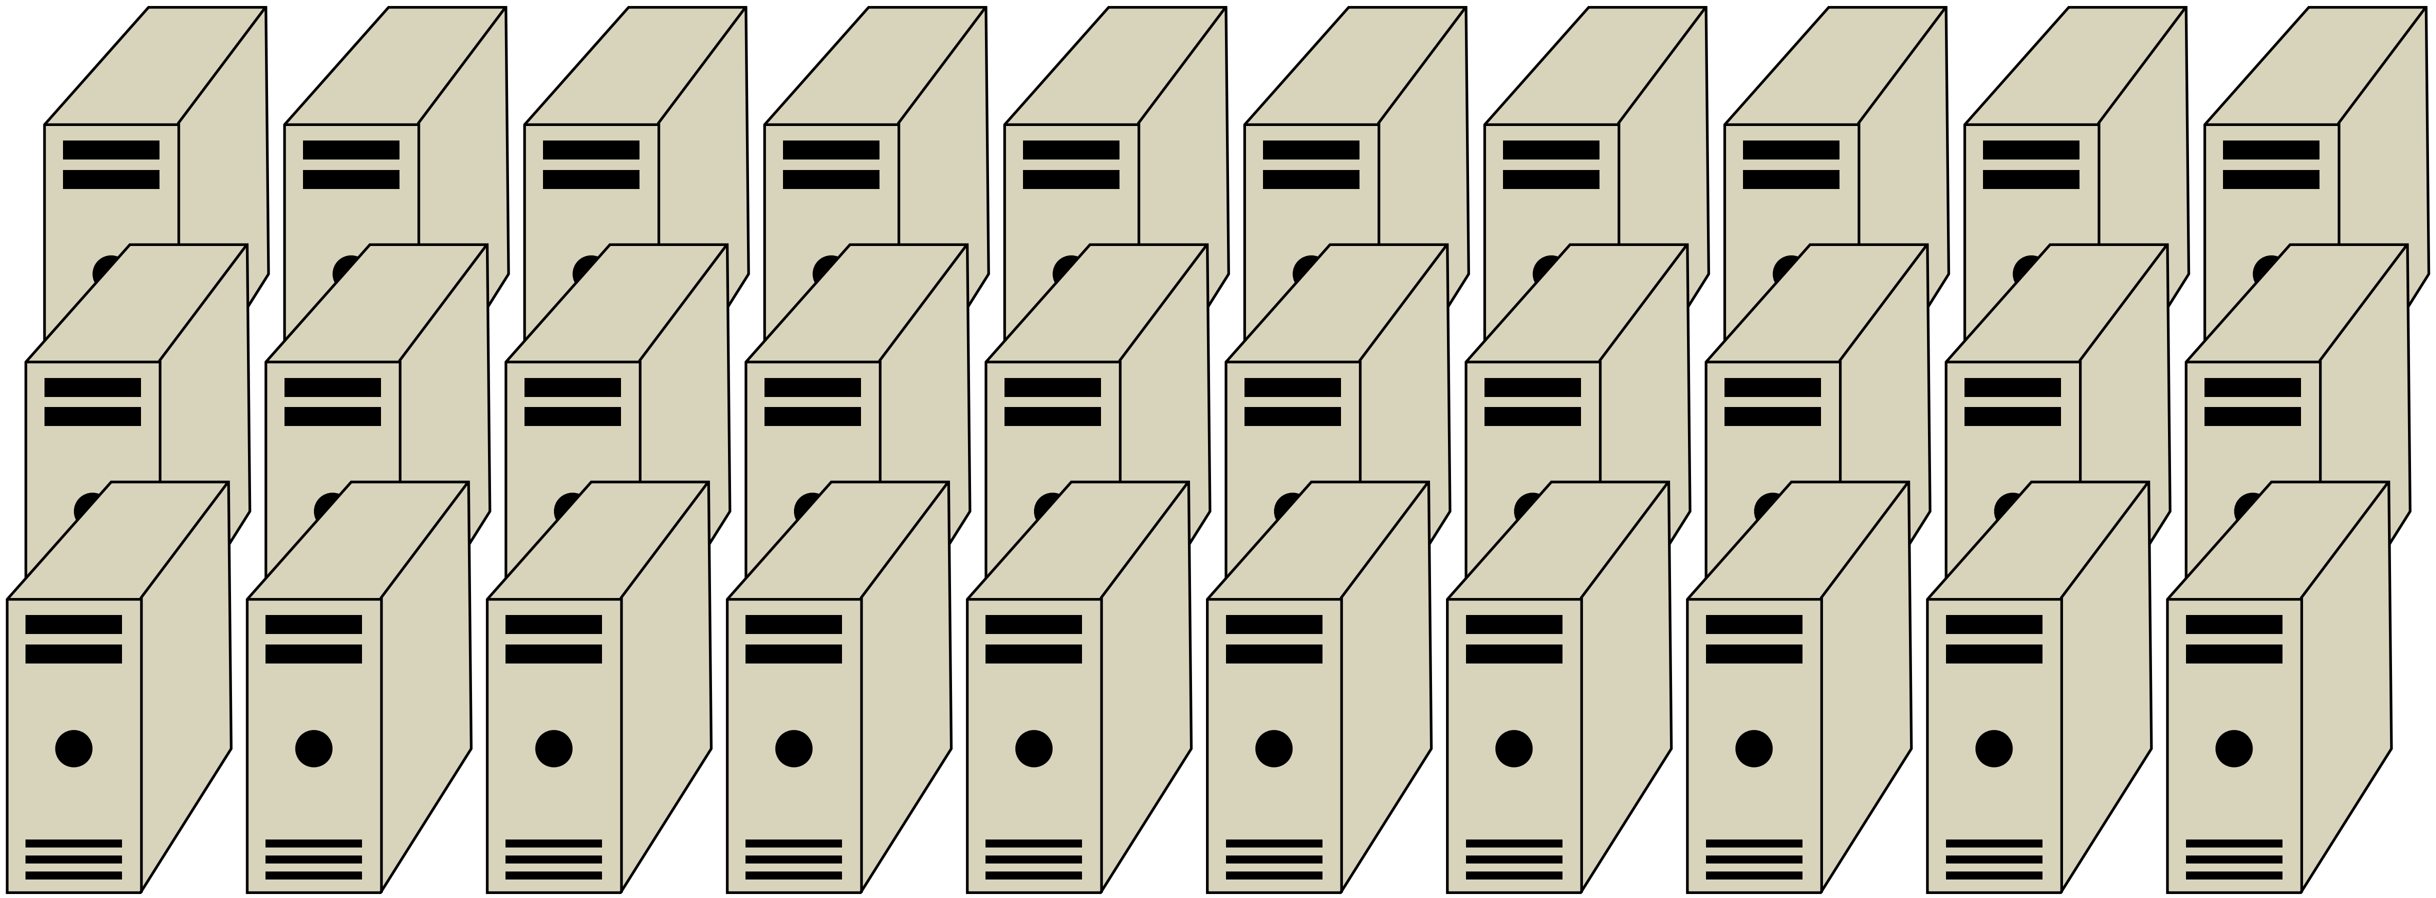

The commercial world has developed techniques for processing PBs of data (Yahoo, Facebook, Amazon).

Scientists are exploring ways of applying these techniques to scientific problems and datasets.

Cloud Computing

- 1000s of commodity computers organized into an on-demand cluster, e.g., Amazon's EC2
- Cheaper than specialized clusters
- Cluster is accessed from anywhere via the internet
- Networking logistics handled automatically
  - Users need very little network computing experience
- Robust to node failures; part of the design
- Nodes easily/rapidly added.

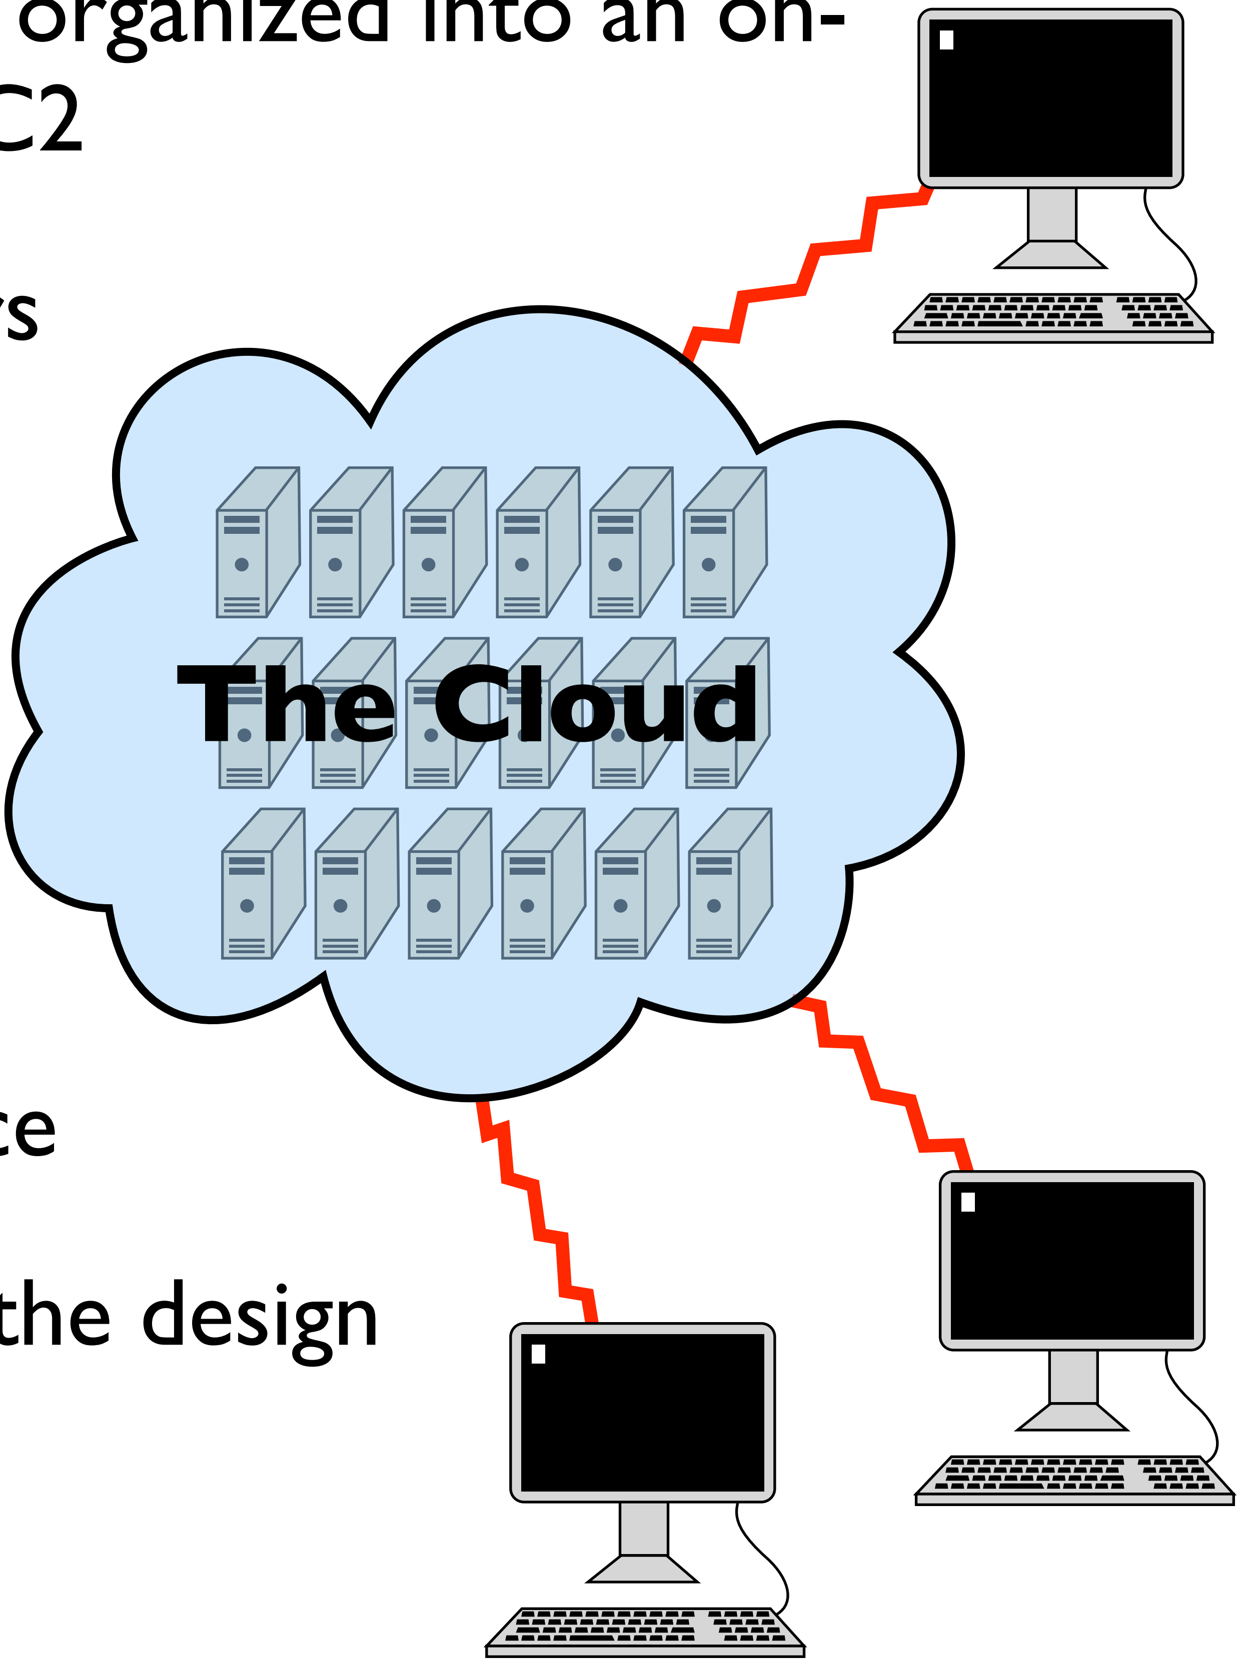

Cloud Computing

We introduce:

MapReduce

(one programming model for cloud computing)

Hadoop

(an implementation of MapReduce)

We will demonstrate ***image coaddition***:

- Given multiple partially overlapping images and a ***query*** (color and sky bounds):
- Find images' intersections with the query bounds.
- Background-subtract, project coordinate system & interpolate (warp), and PSF*-match intersections.
- Weight, stack, and mosaic into a final product.

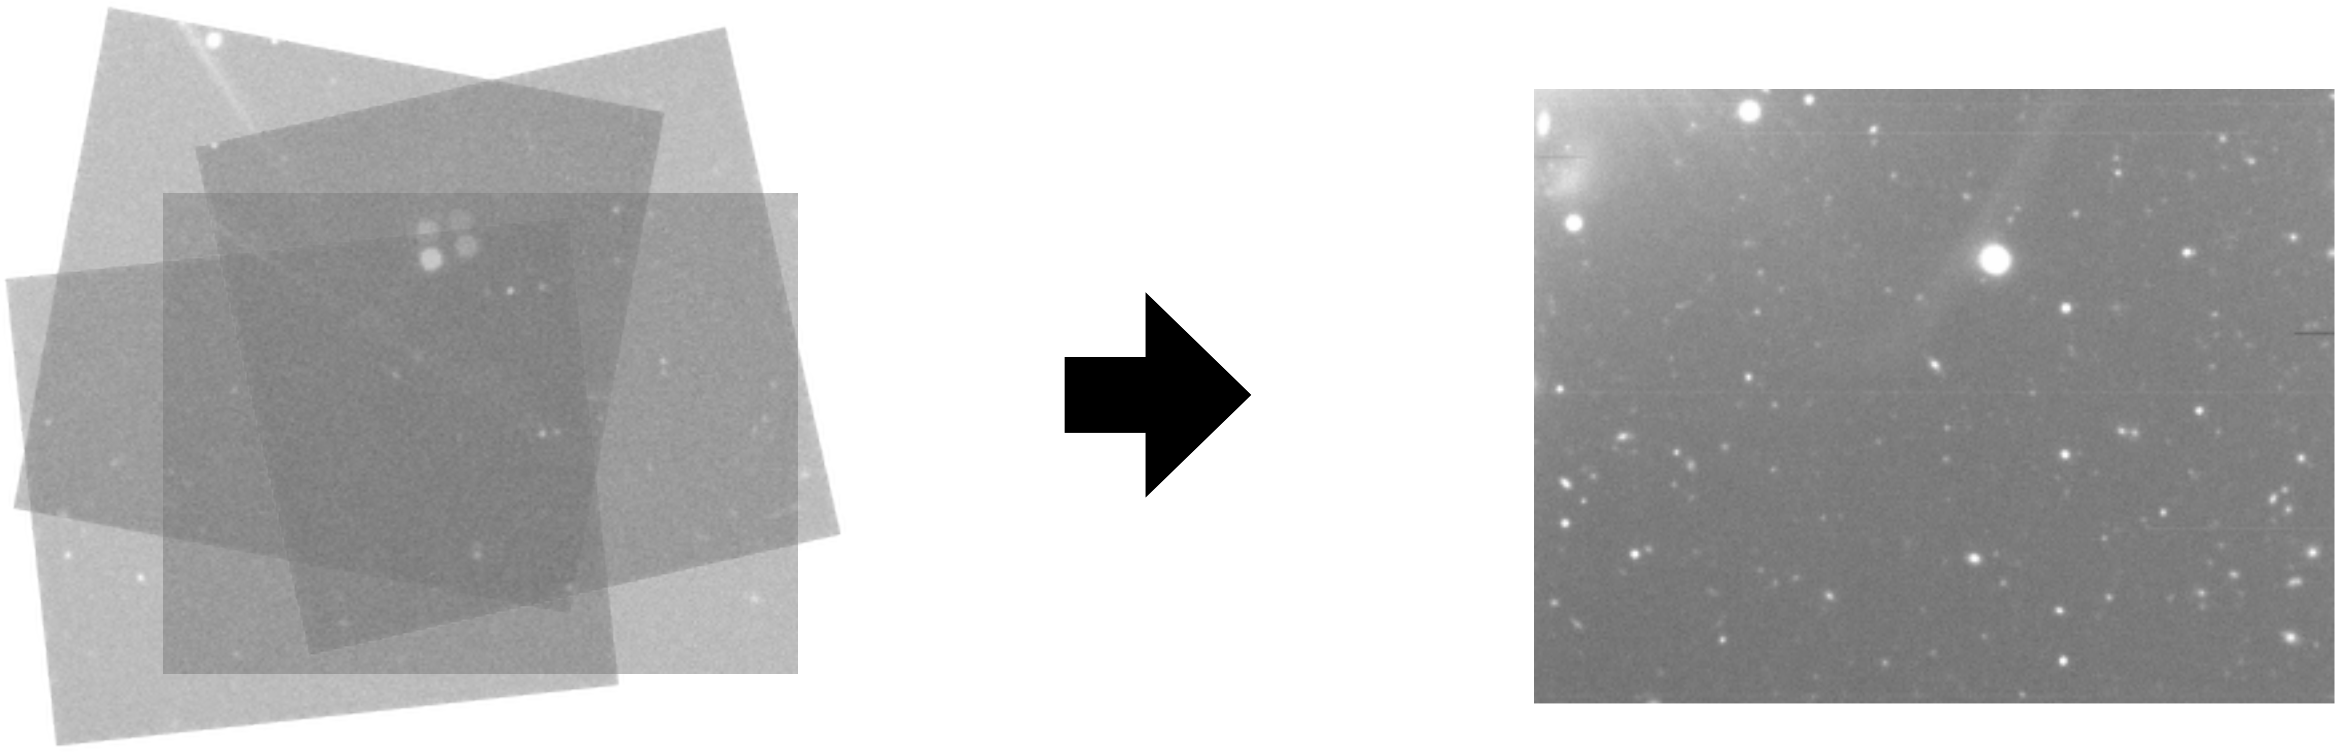

* Point-spread function

*SDSS**

Camera has 30 CCDs:
- 5 bandpass filters
- 6 abutting strips of sky
- 2048x1489 pixels per CCD (~6MB uncompressed FITS)

Stripe 82 dataset:
30 TBs, 4 million images

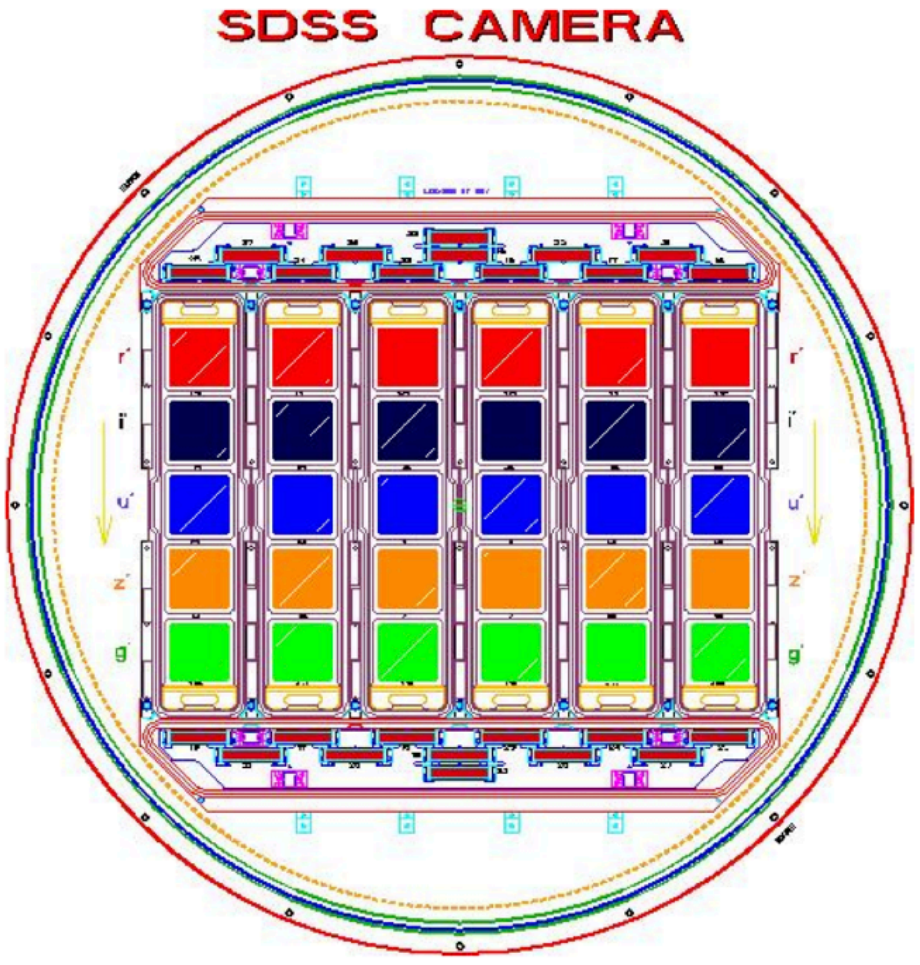

* Sloan Digital Sky Survey

MapReduce

A massively parallel database-processing framework

In one sense:
A parallel database

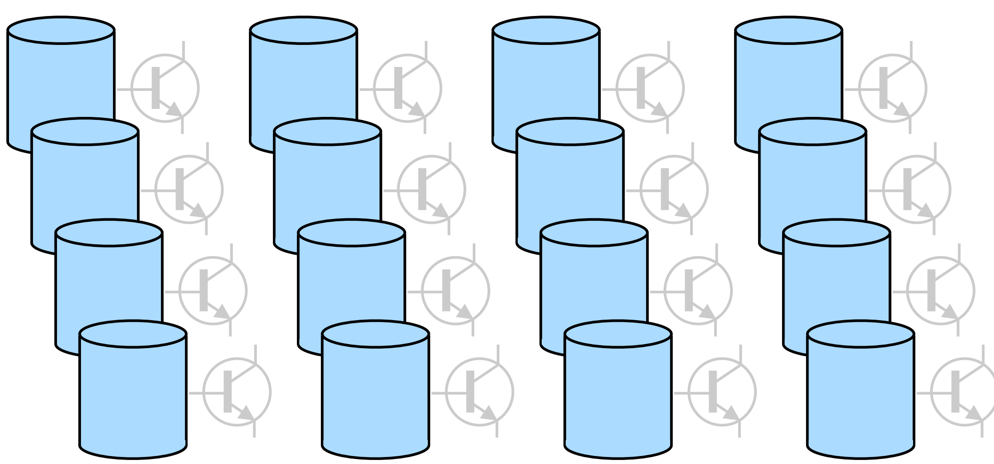

In another sense:
A parallel computing cluster

It's both!

MapReduce

1. ***Mappers*** process local data to an intermediate state.
2. Mapper outputs are ***shuffled*** to reducers.
3. ***Reducers*** further process the data, producing final output.

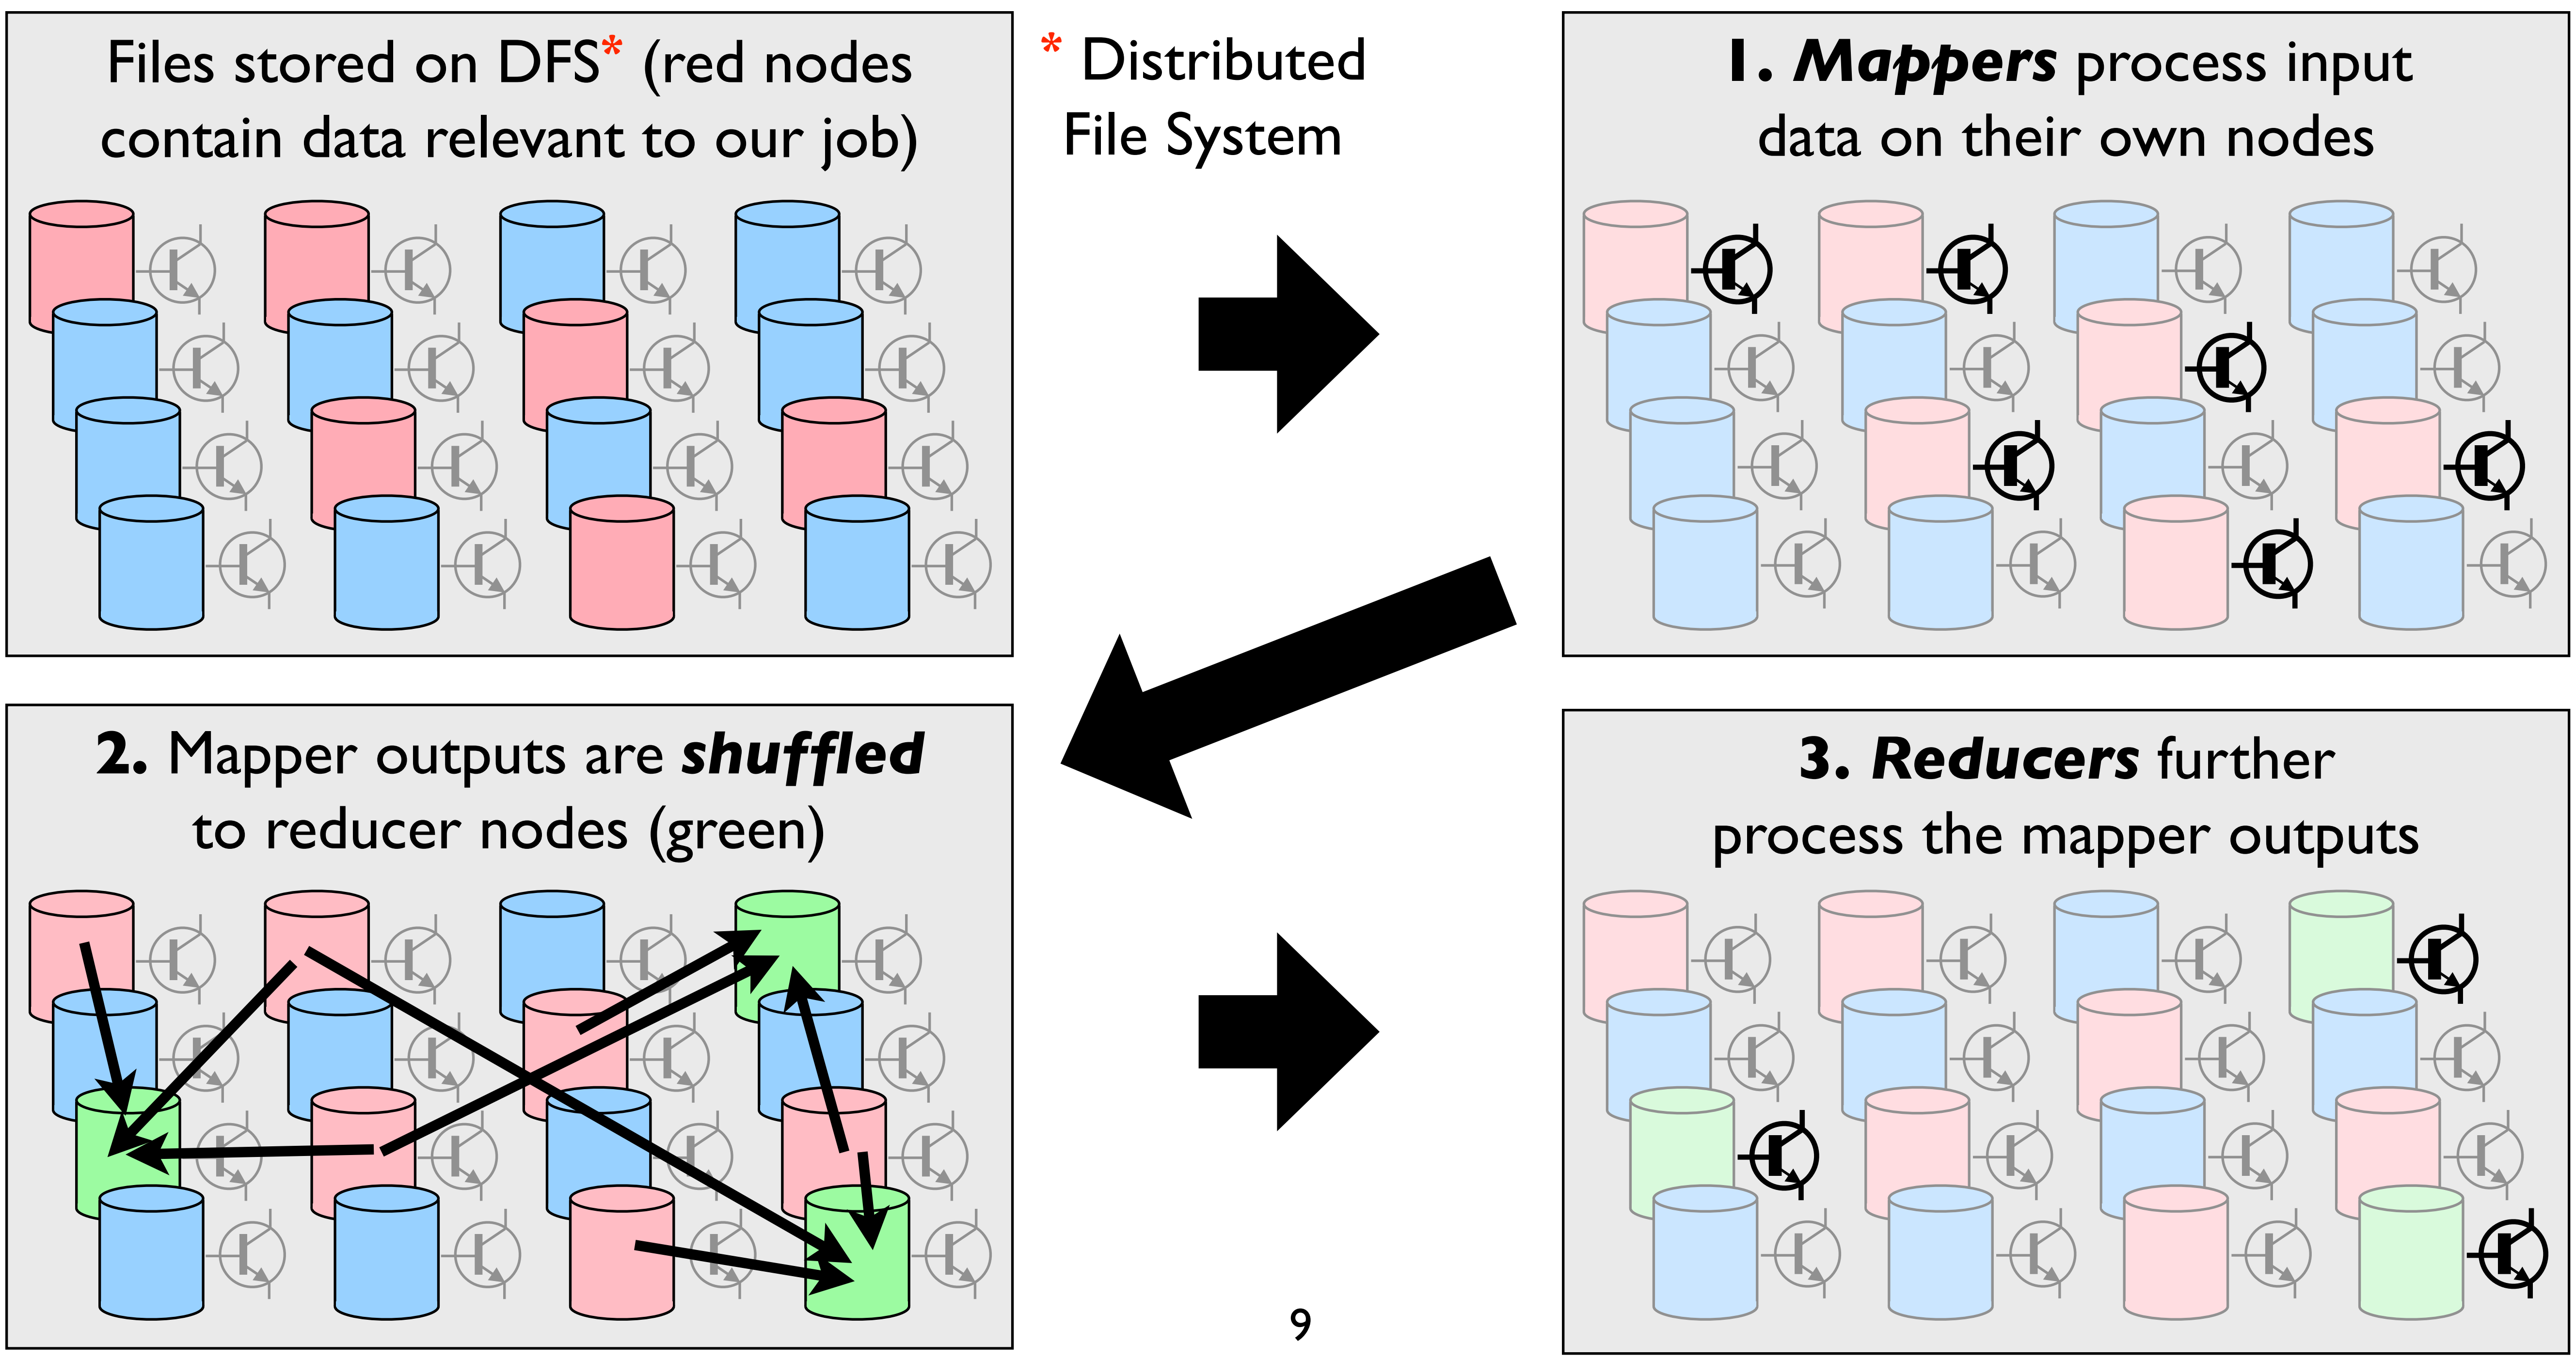

Apache Hadoop
An implementation of MapReduce

- Open source, largely contributed by Yahoo
- Implemented in Java
- Programmed in Java
- Widely used in industry (Yahoo, Facebook, Amazon)
- Active user community (good support base)

Hadoop is implemented and programmed in Java.

However, we want to use a powerful (compiled) C++ image processing library.

JNI* facilitates the coupling between the two components.

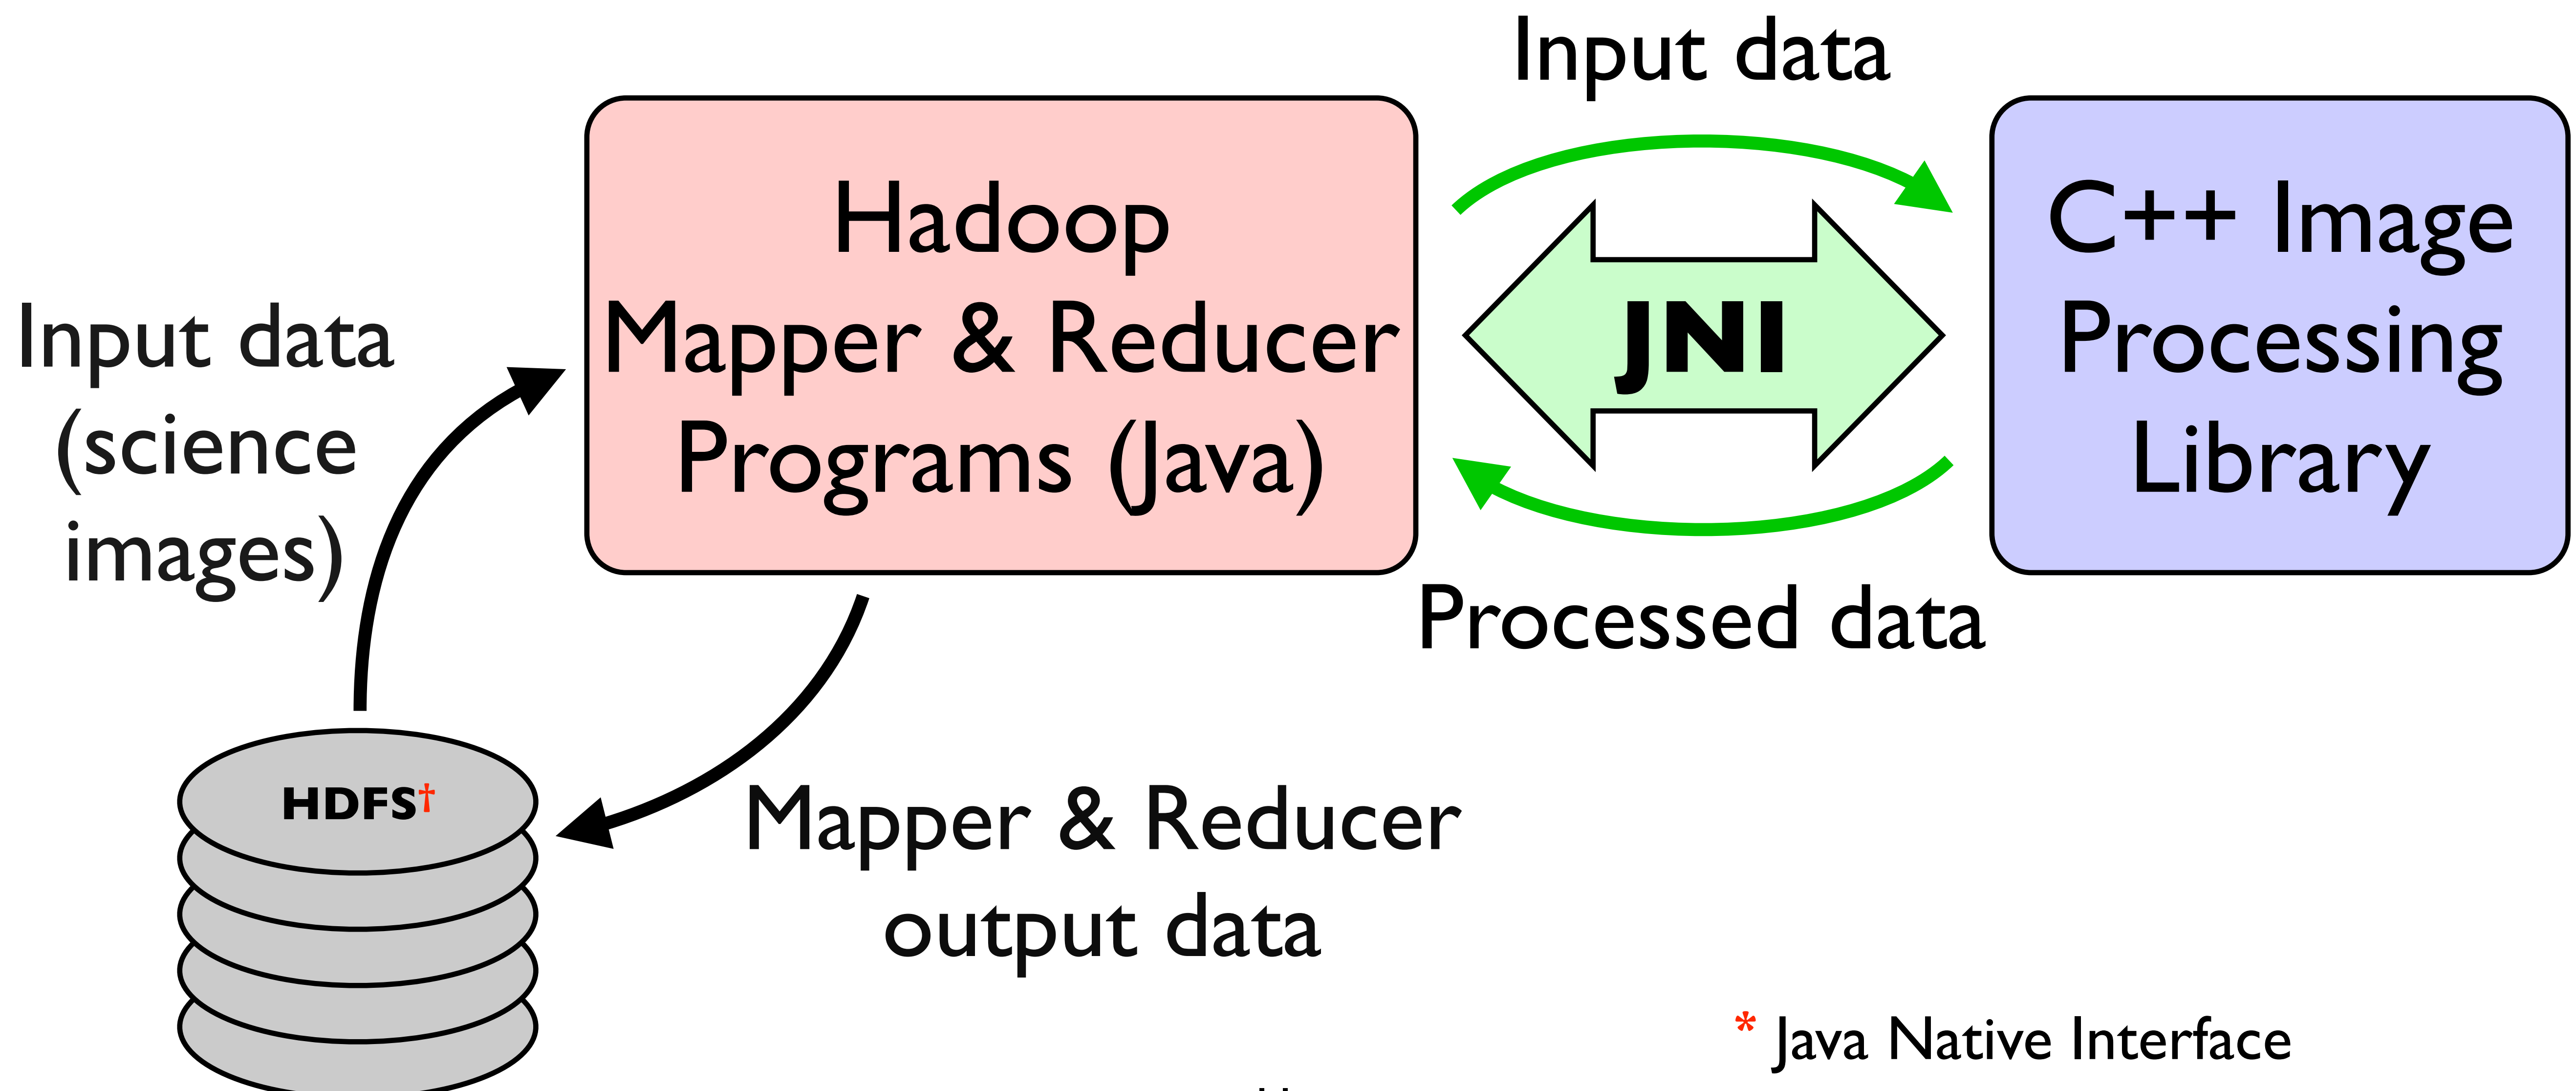

* Java Native Interface
† Hadoop Distributed File System

Image Coaddition with SQL and Hadoop

We only need a tiny fraction of the total images from the database to process a given query (color and sky bounds).

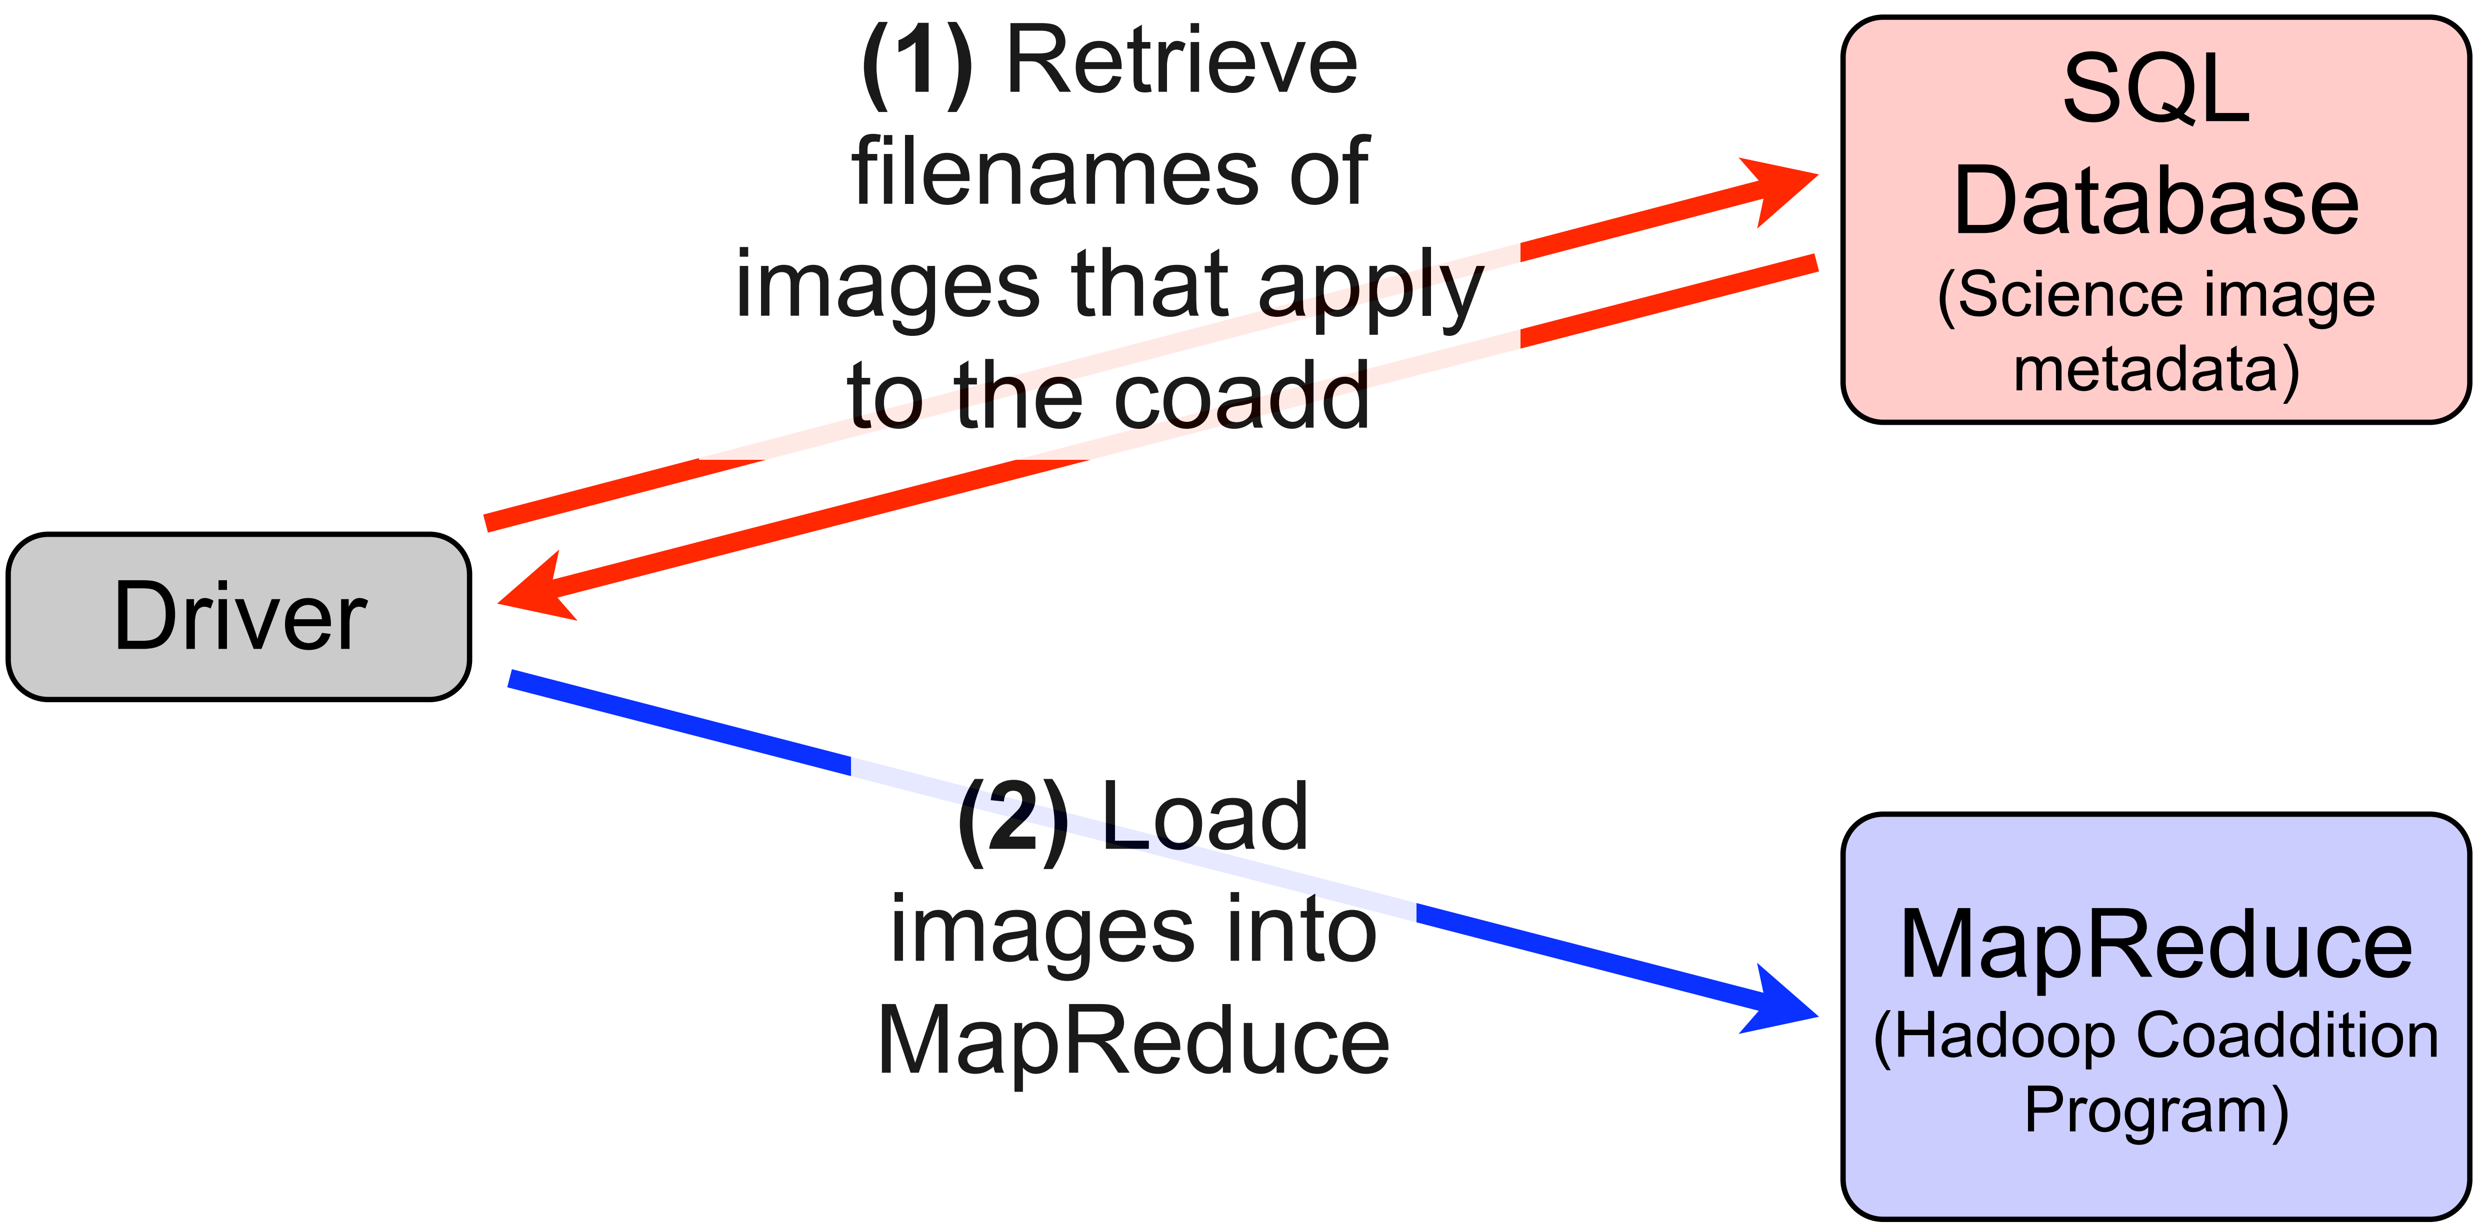

Image Coaddition in Hadoop

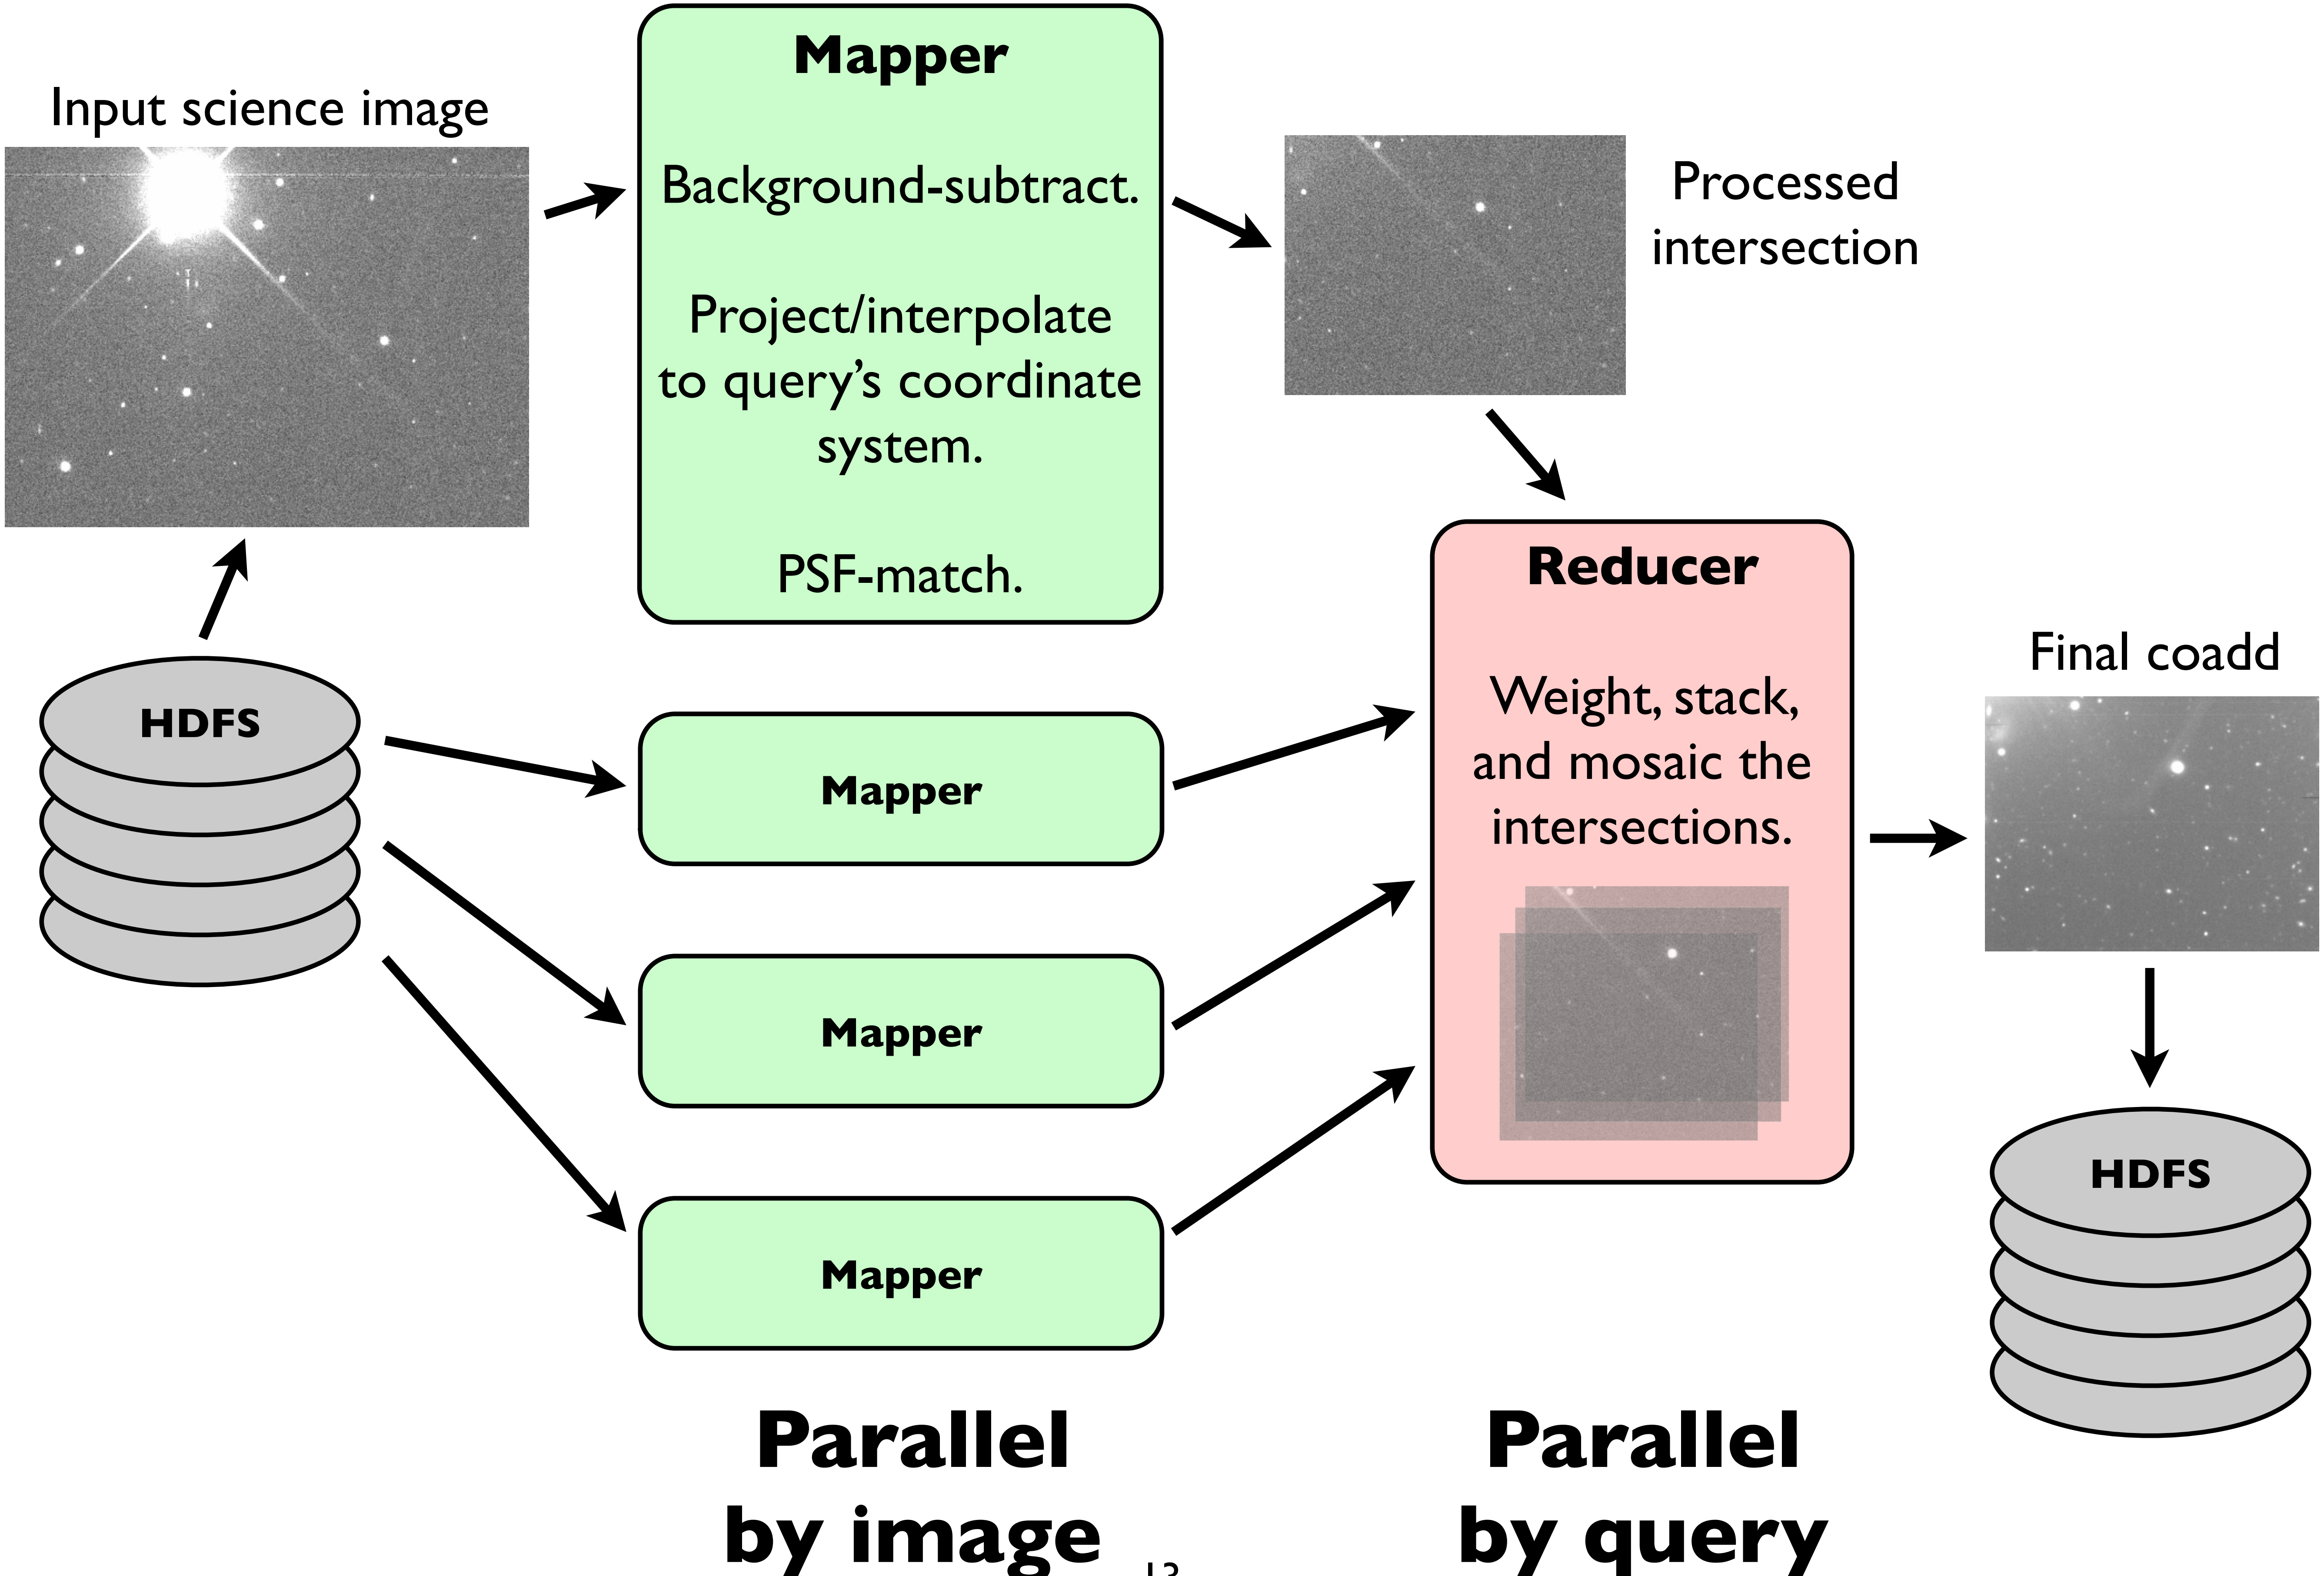

Example

SDSS 2570-r6-199

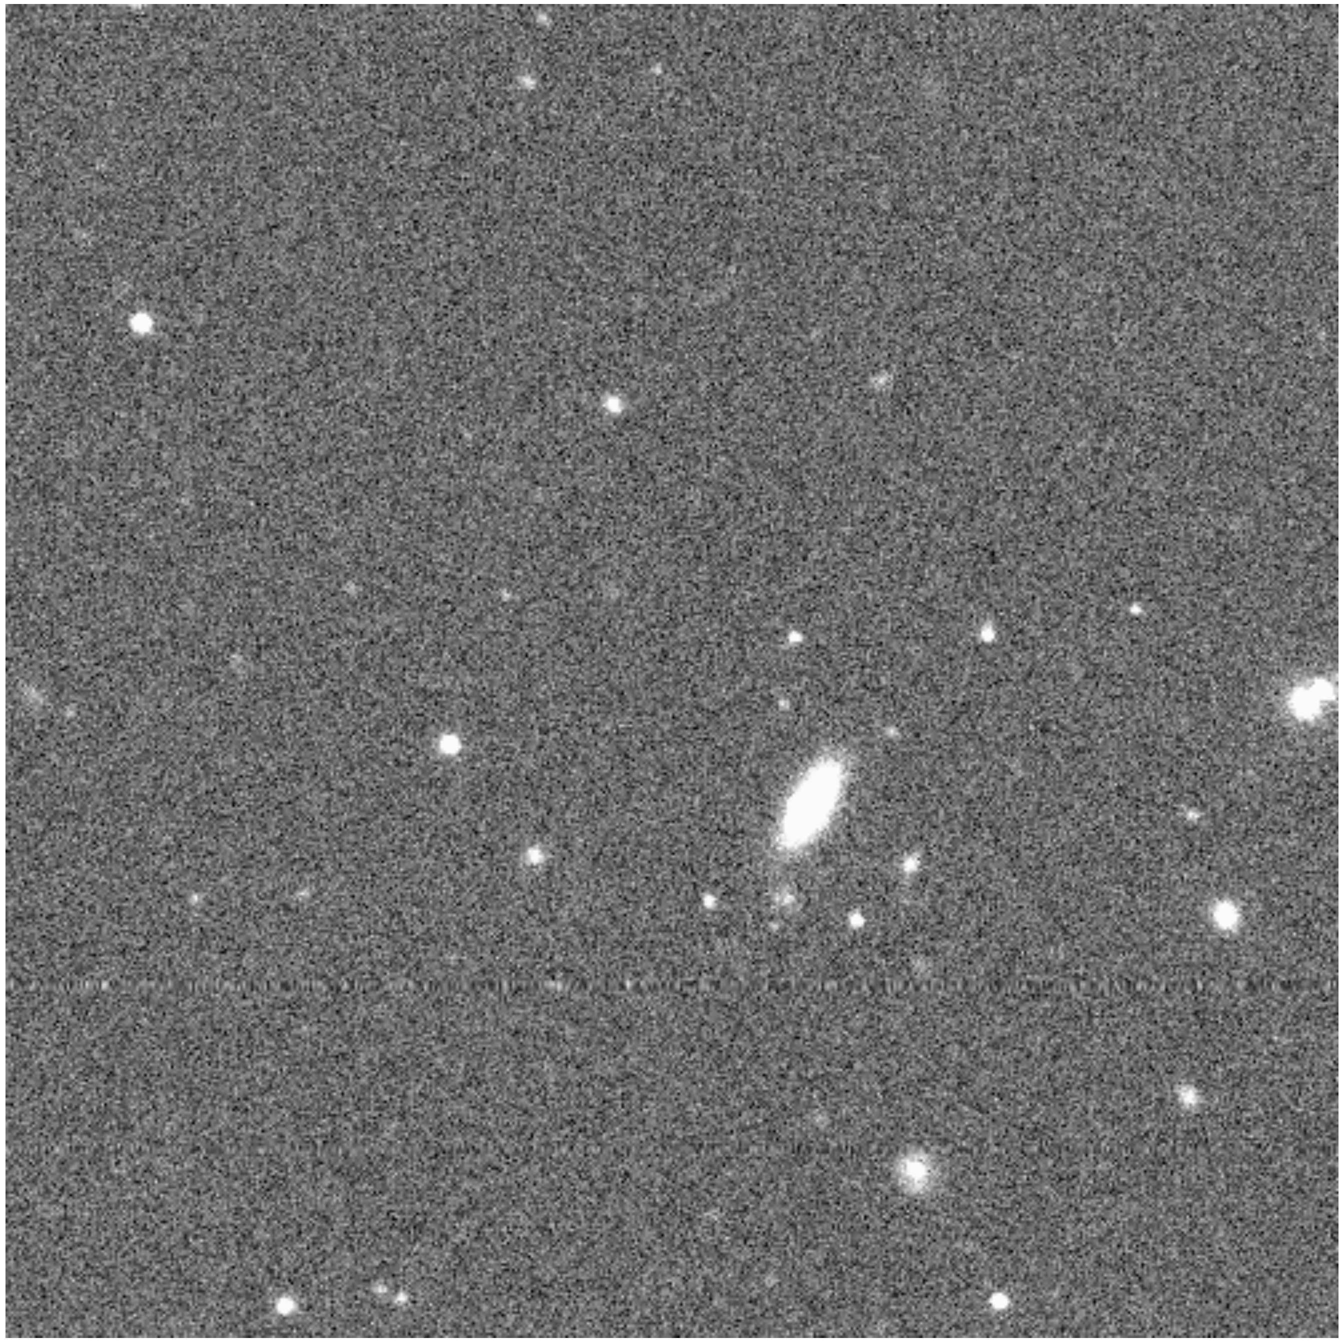

Coadd of 96 images*

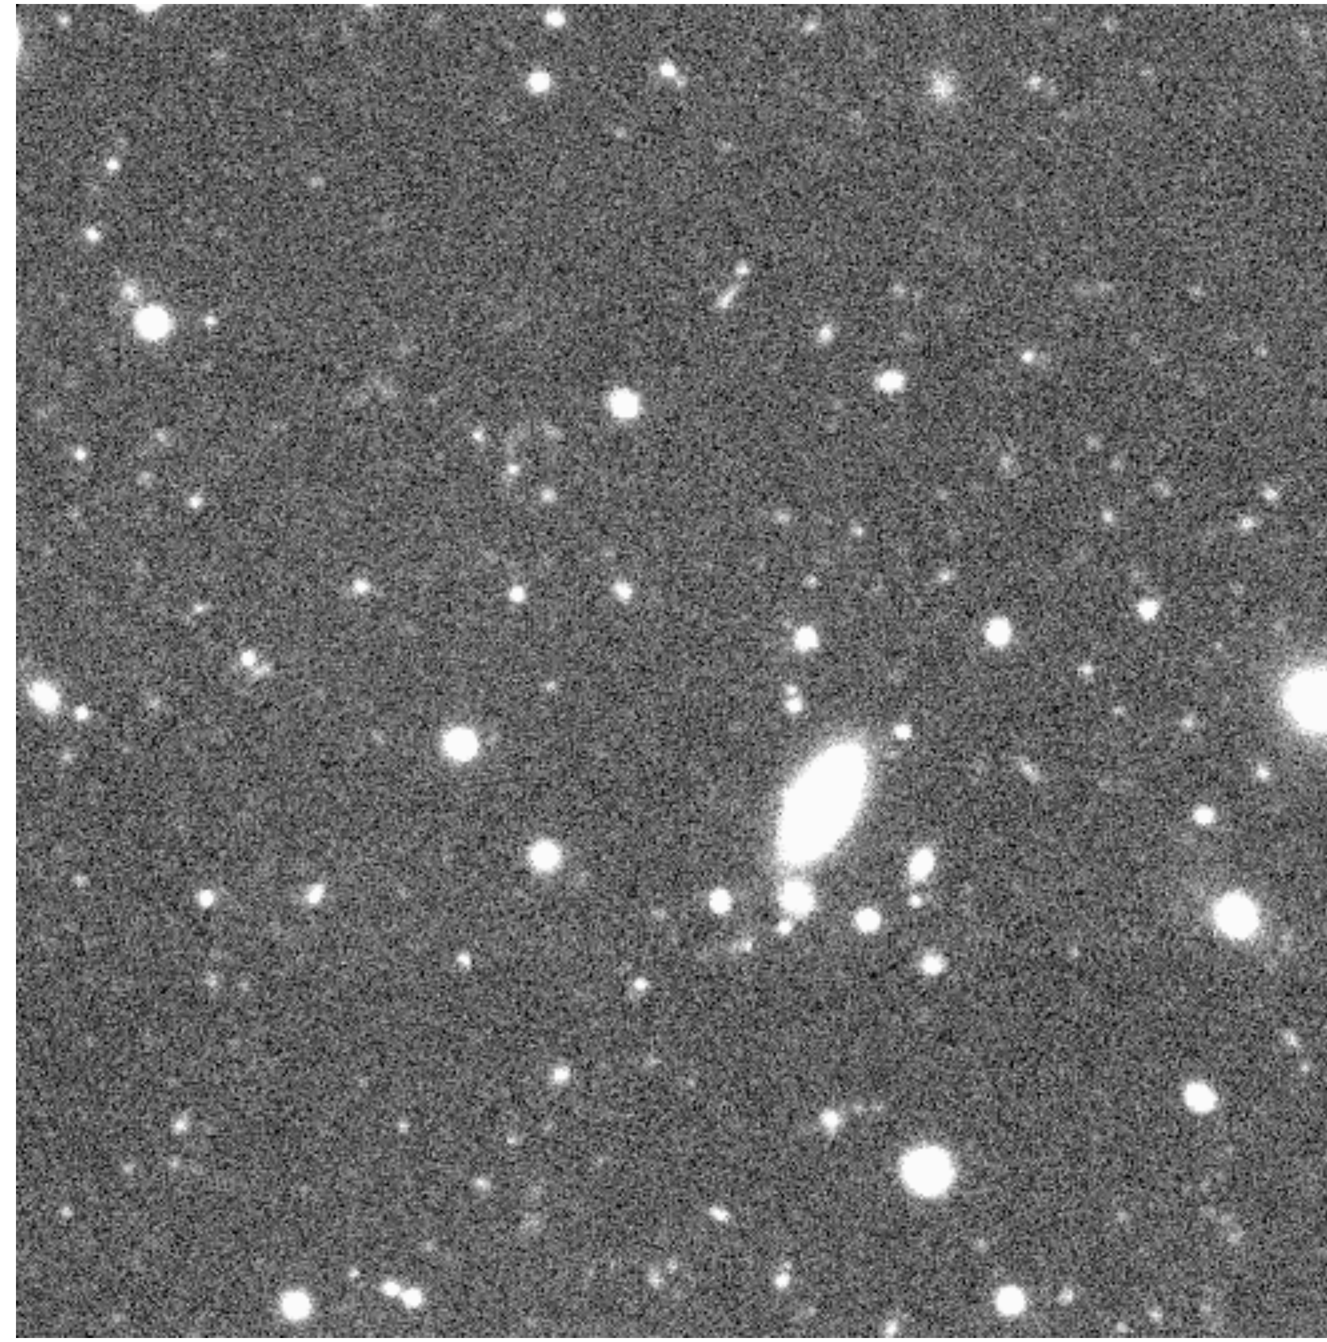

Expected improved limiting magnitude =
$-2.5\log(\sqrt{96}) \approx -2.5$ mags

*Coverage is not necessarily 96 at any given pixel

Limiting Magnitude Comparison

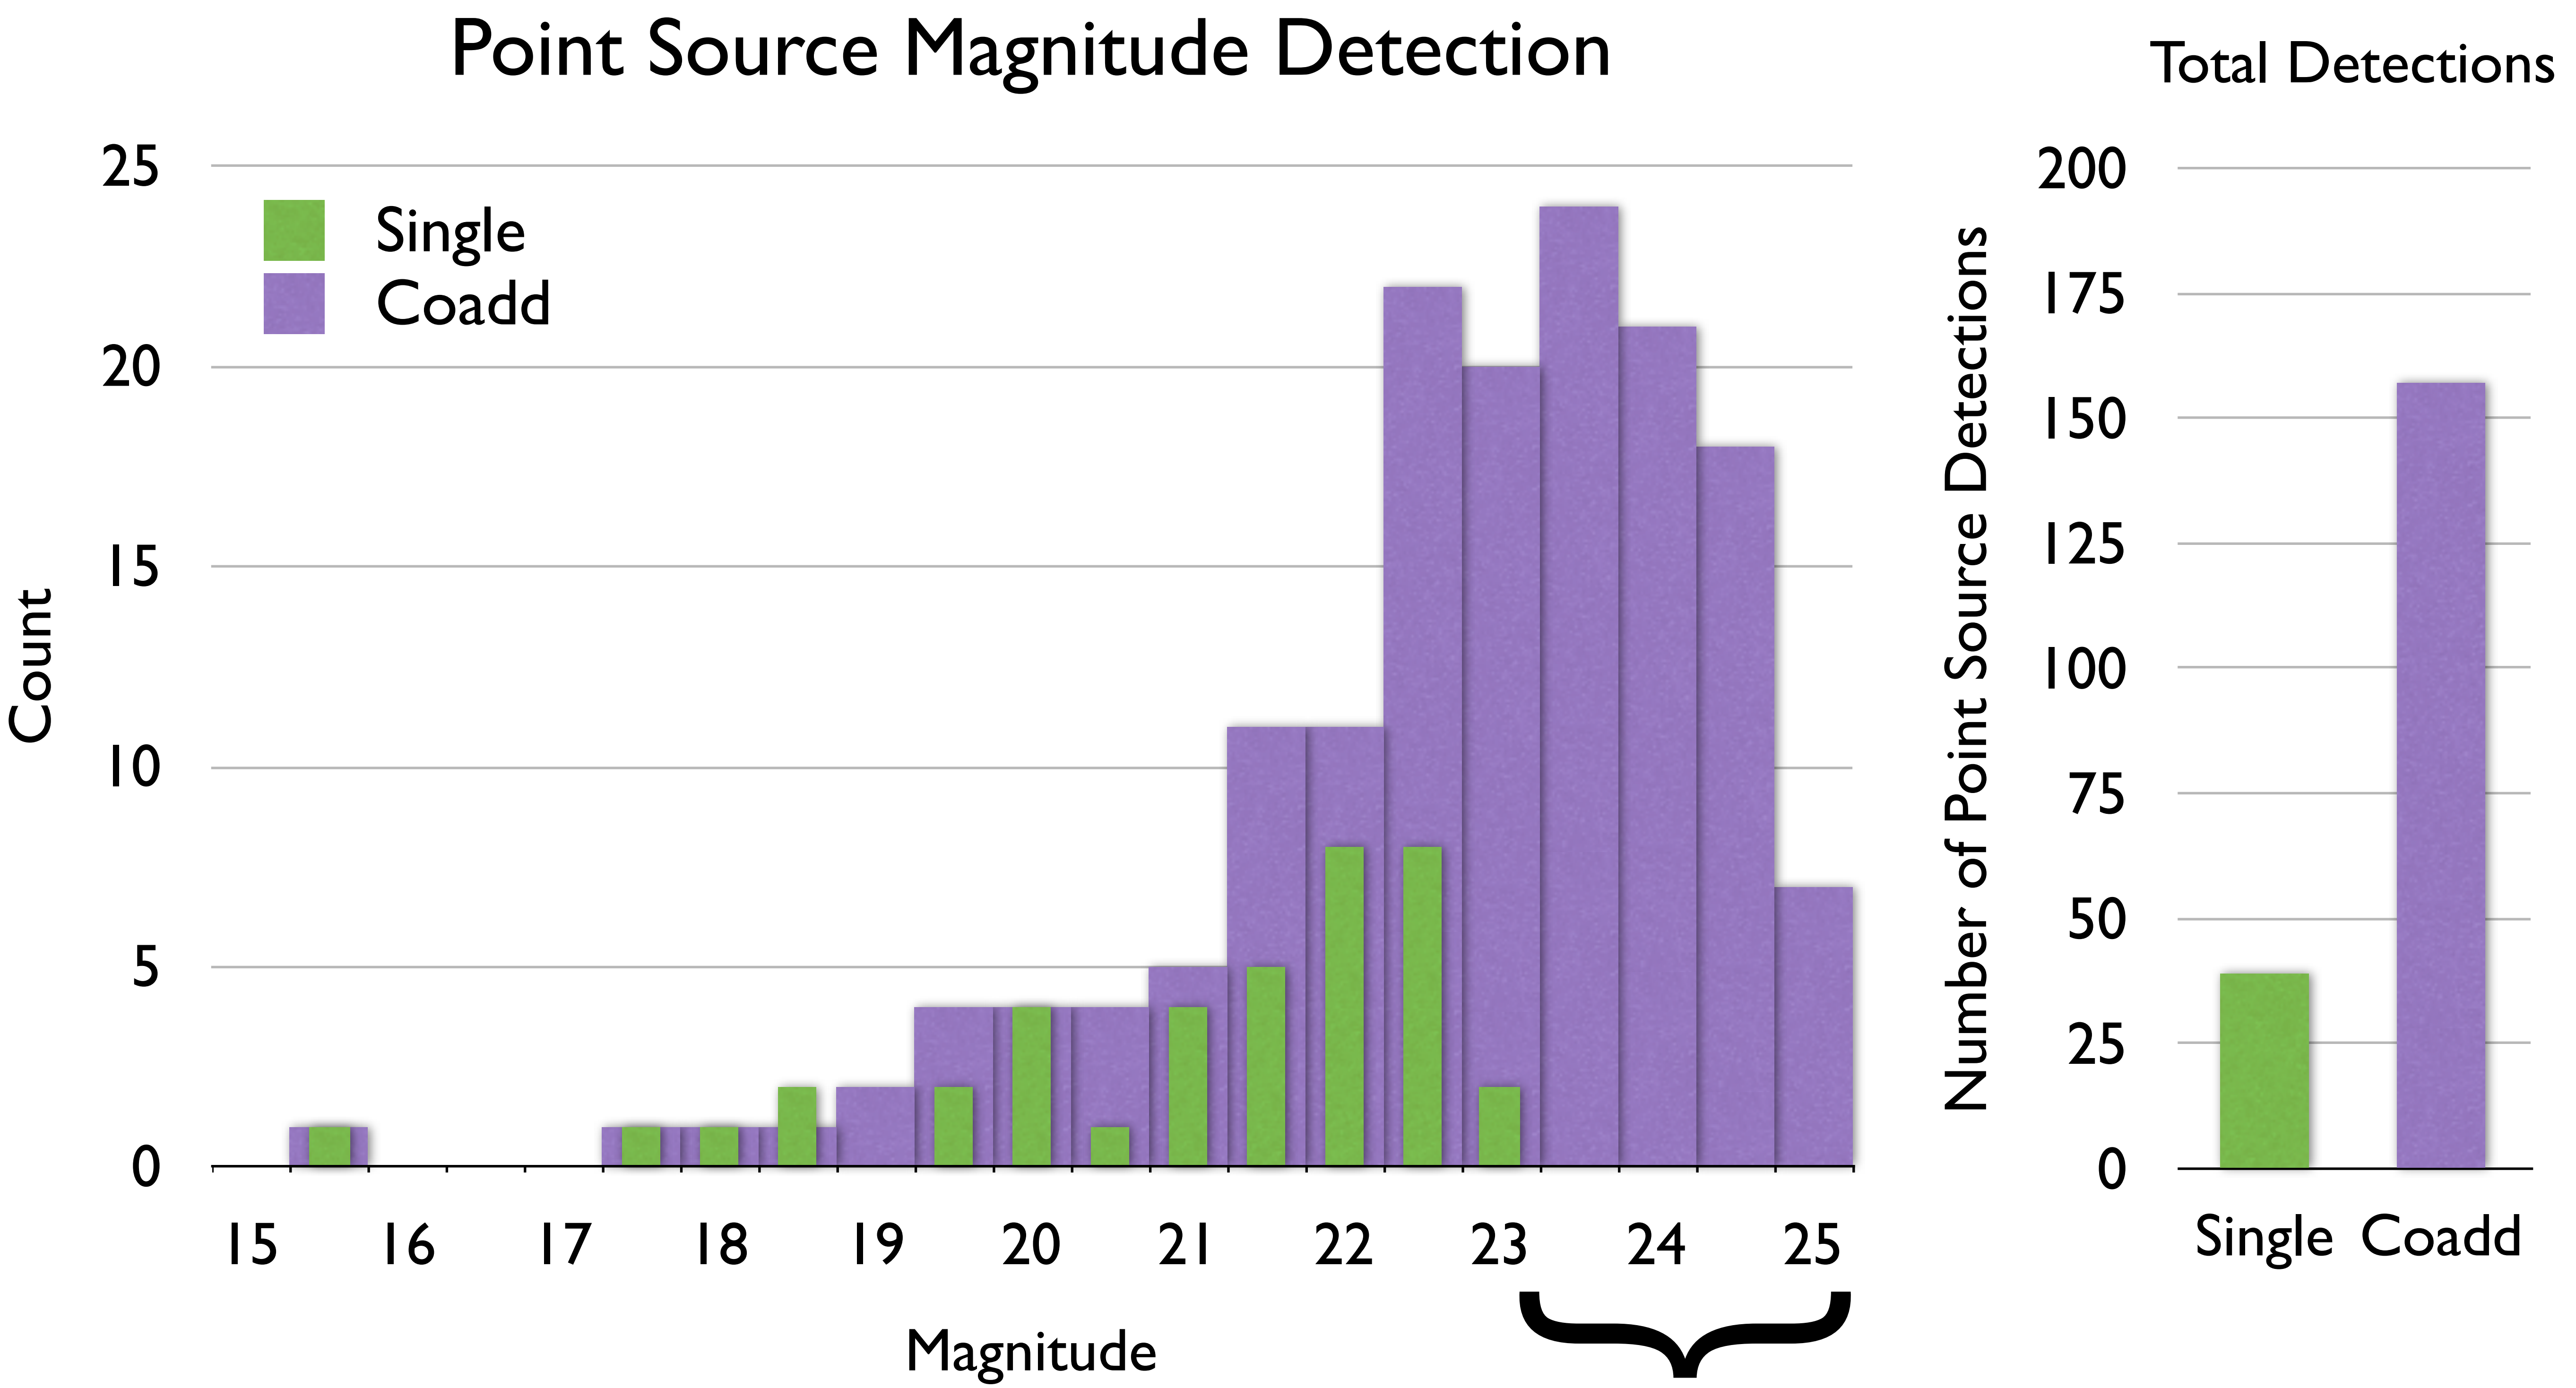

We gained ~2 mags in point source detection depth*

* As expected for a 96x stack (see previous slide)

CluE* Cluster Configuration

- ~700 nodes:
  - 2 processors 2.8Ghz Xeon (dual core) (4 cores per node)
  - 8GB RAM (2GB per task)
  - 2 disks 400GBs (560TBs on cluster)
- ~1400 mapper slots, ~1400 reducer slots

* NSF ***Clu**ster **E**xploratory* Grant,
cluster maintained by Google/IBM.

Running Time for the Coadd Shown in this Talk:

- 170 images returned by SQL (sent to mappers)
- 96 intersections coadded by reducer (many mappers fail to find good PSF-matching candidates, *i.e.*, high-quality stars)
- SQL query: 2 mins
- Mappers: 29 mins (8 mins w/o retries)
- Reducer: 1.5 mins
- **Total:** **34 mins** (13 mins w/o retries)

Conclusions

- ***Stored:***
  - SDSS Stripe 82 on a Hadoop cluster (HDFS) (30 TBs, 4 million images)
  - Color/sky-bounds metadata in a SQL database.

- ***Generated high-quality coadds:***
  - Background-subtraction
  - Coordinate system projection/interpolation
  - PSF-matching
  - Weighted stacking
  - ***Time:*** 15 to 60 minutes per 500x500px coadd

Future Work

- ***Improve the overall algorithm:***
  - Parallelize reducer
  - Better memory management
  - Simultaneous queries
  - Improved background-subtraction, PSF-matching, etc.
- ***Time-bounded queries & followup analysis:***
  - Detection of moving/transient objects
  - Automated object detection and classification.
- ***More user-friendly interface:***
  - Higher-level languages that wrap Hadoop (Pig, Hive)
  - GUI front-end (web-interface).

Questions?

Keith Wiley
kbwiley@astro.washington.edu