Analyzing astronomical data with Apache Spark

Christian Arnault
CNRS, Laboratoire de l'Accélérateur Linéaire

XXIst century astronomy?

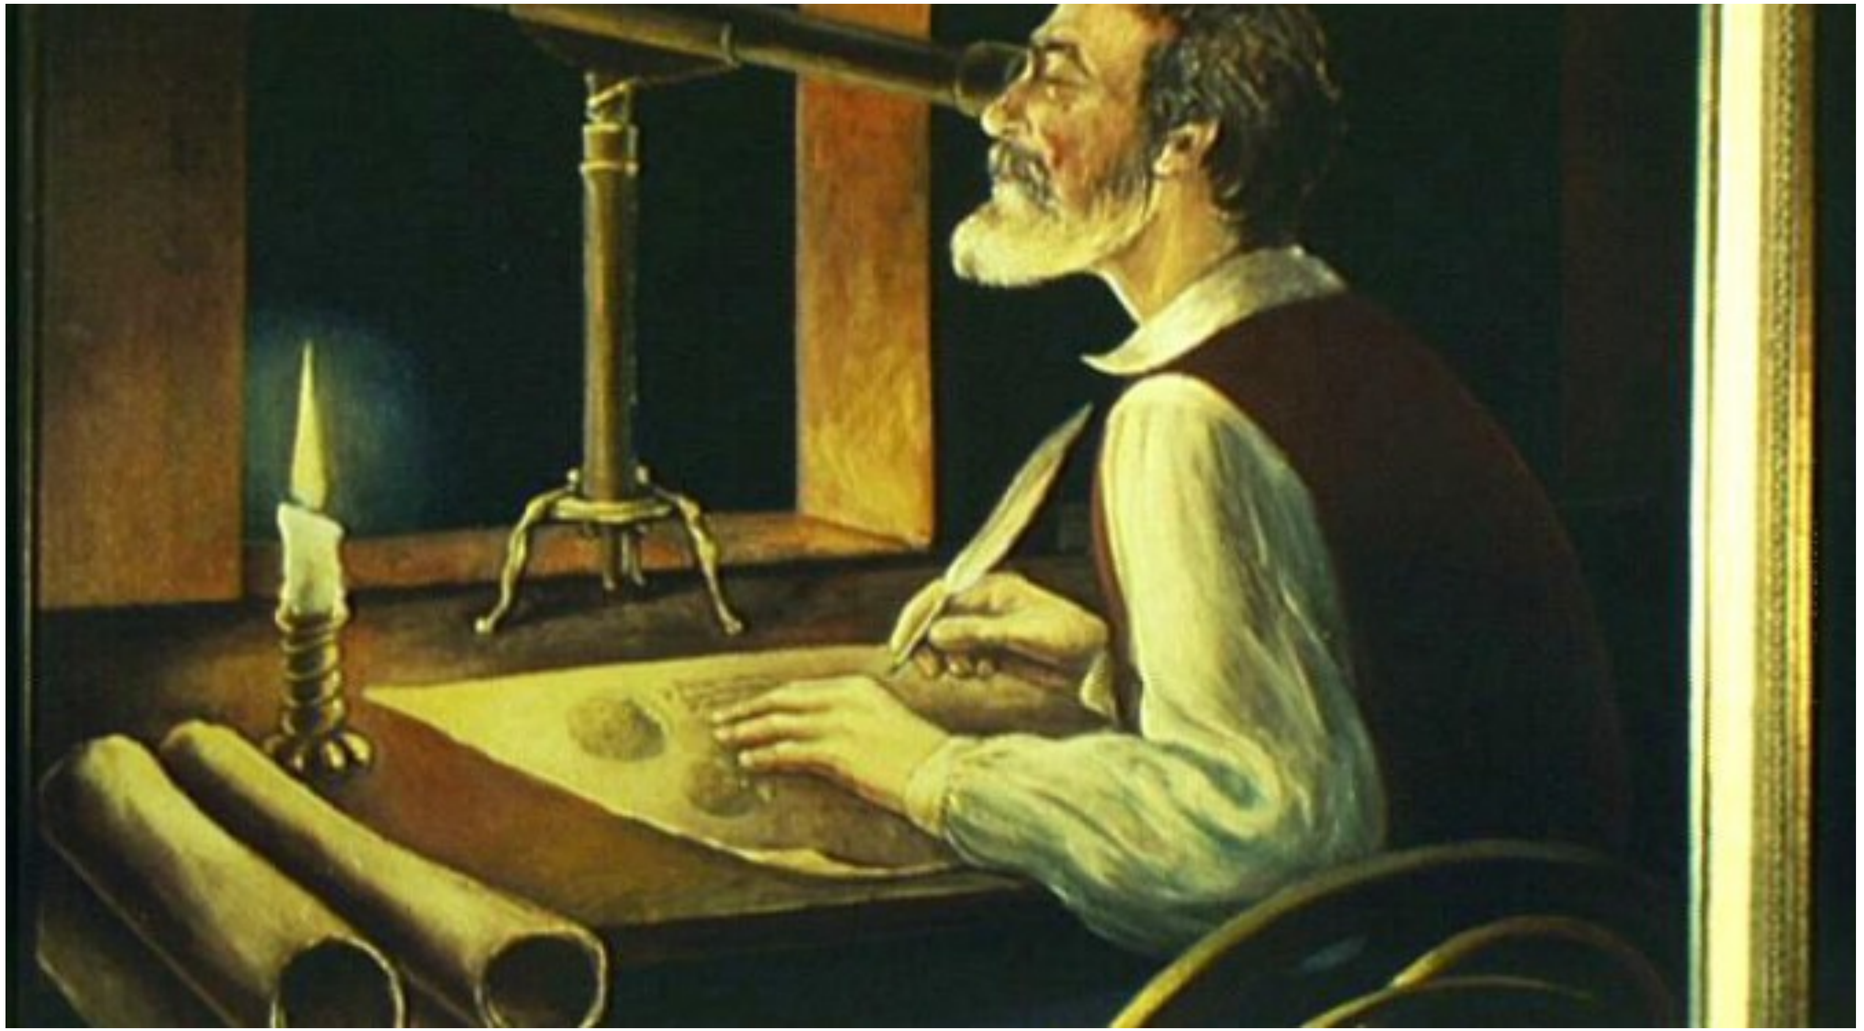

Catalog objects evolution (million objects)

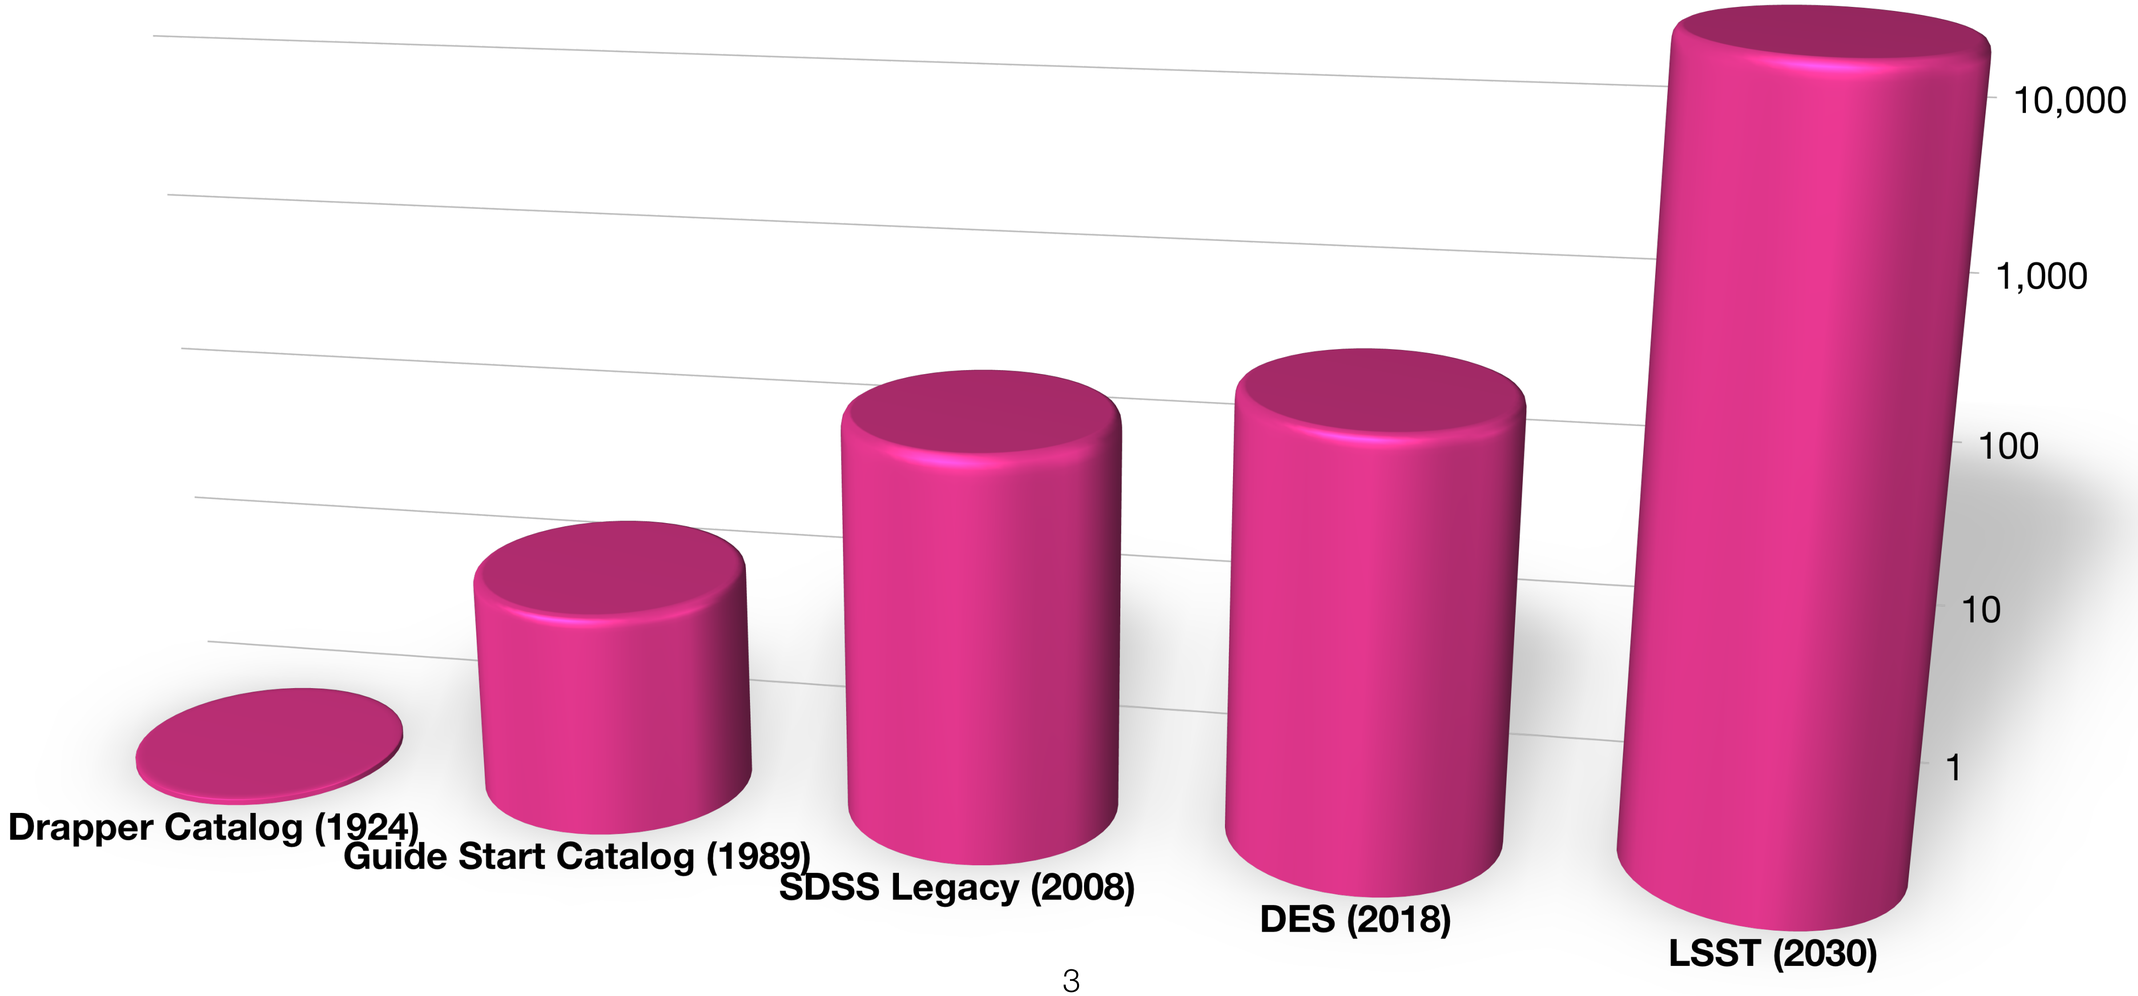

Catalog objects evolution (million objects)

Data being better vs data being *different*

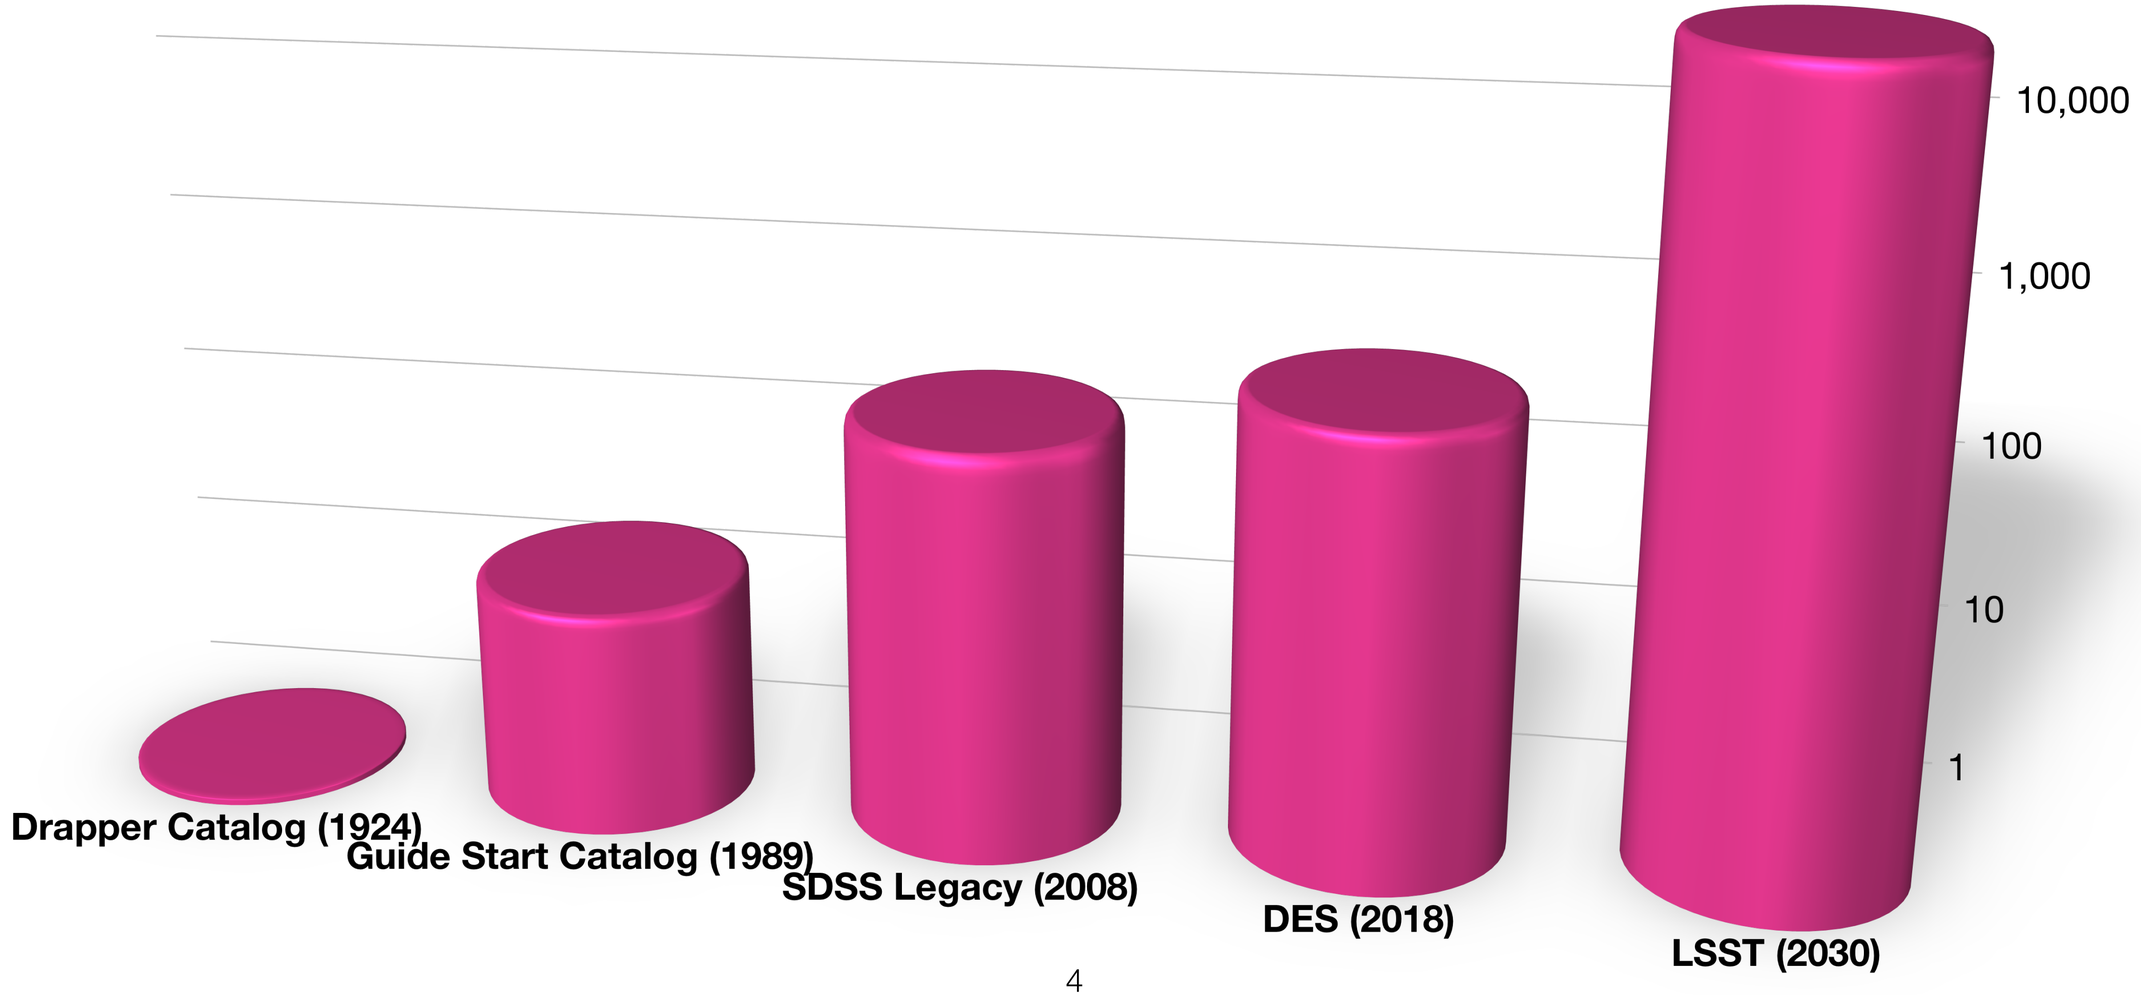

Large Synoptic Survey Telescope

- Non-profit corporation
- Chile (Cerro Pachón)
- US led, international collaboration
- 2020 - 2030

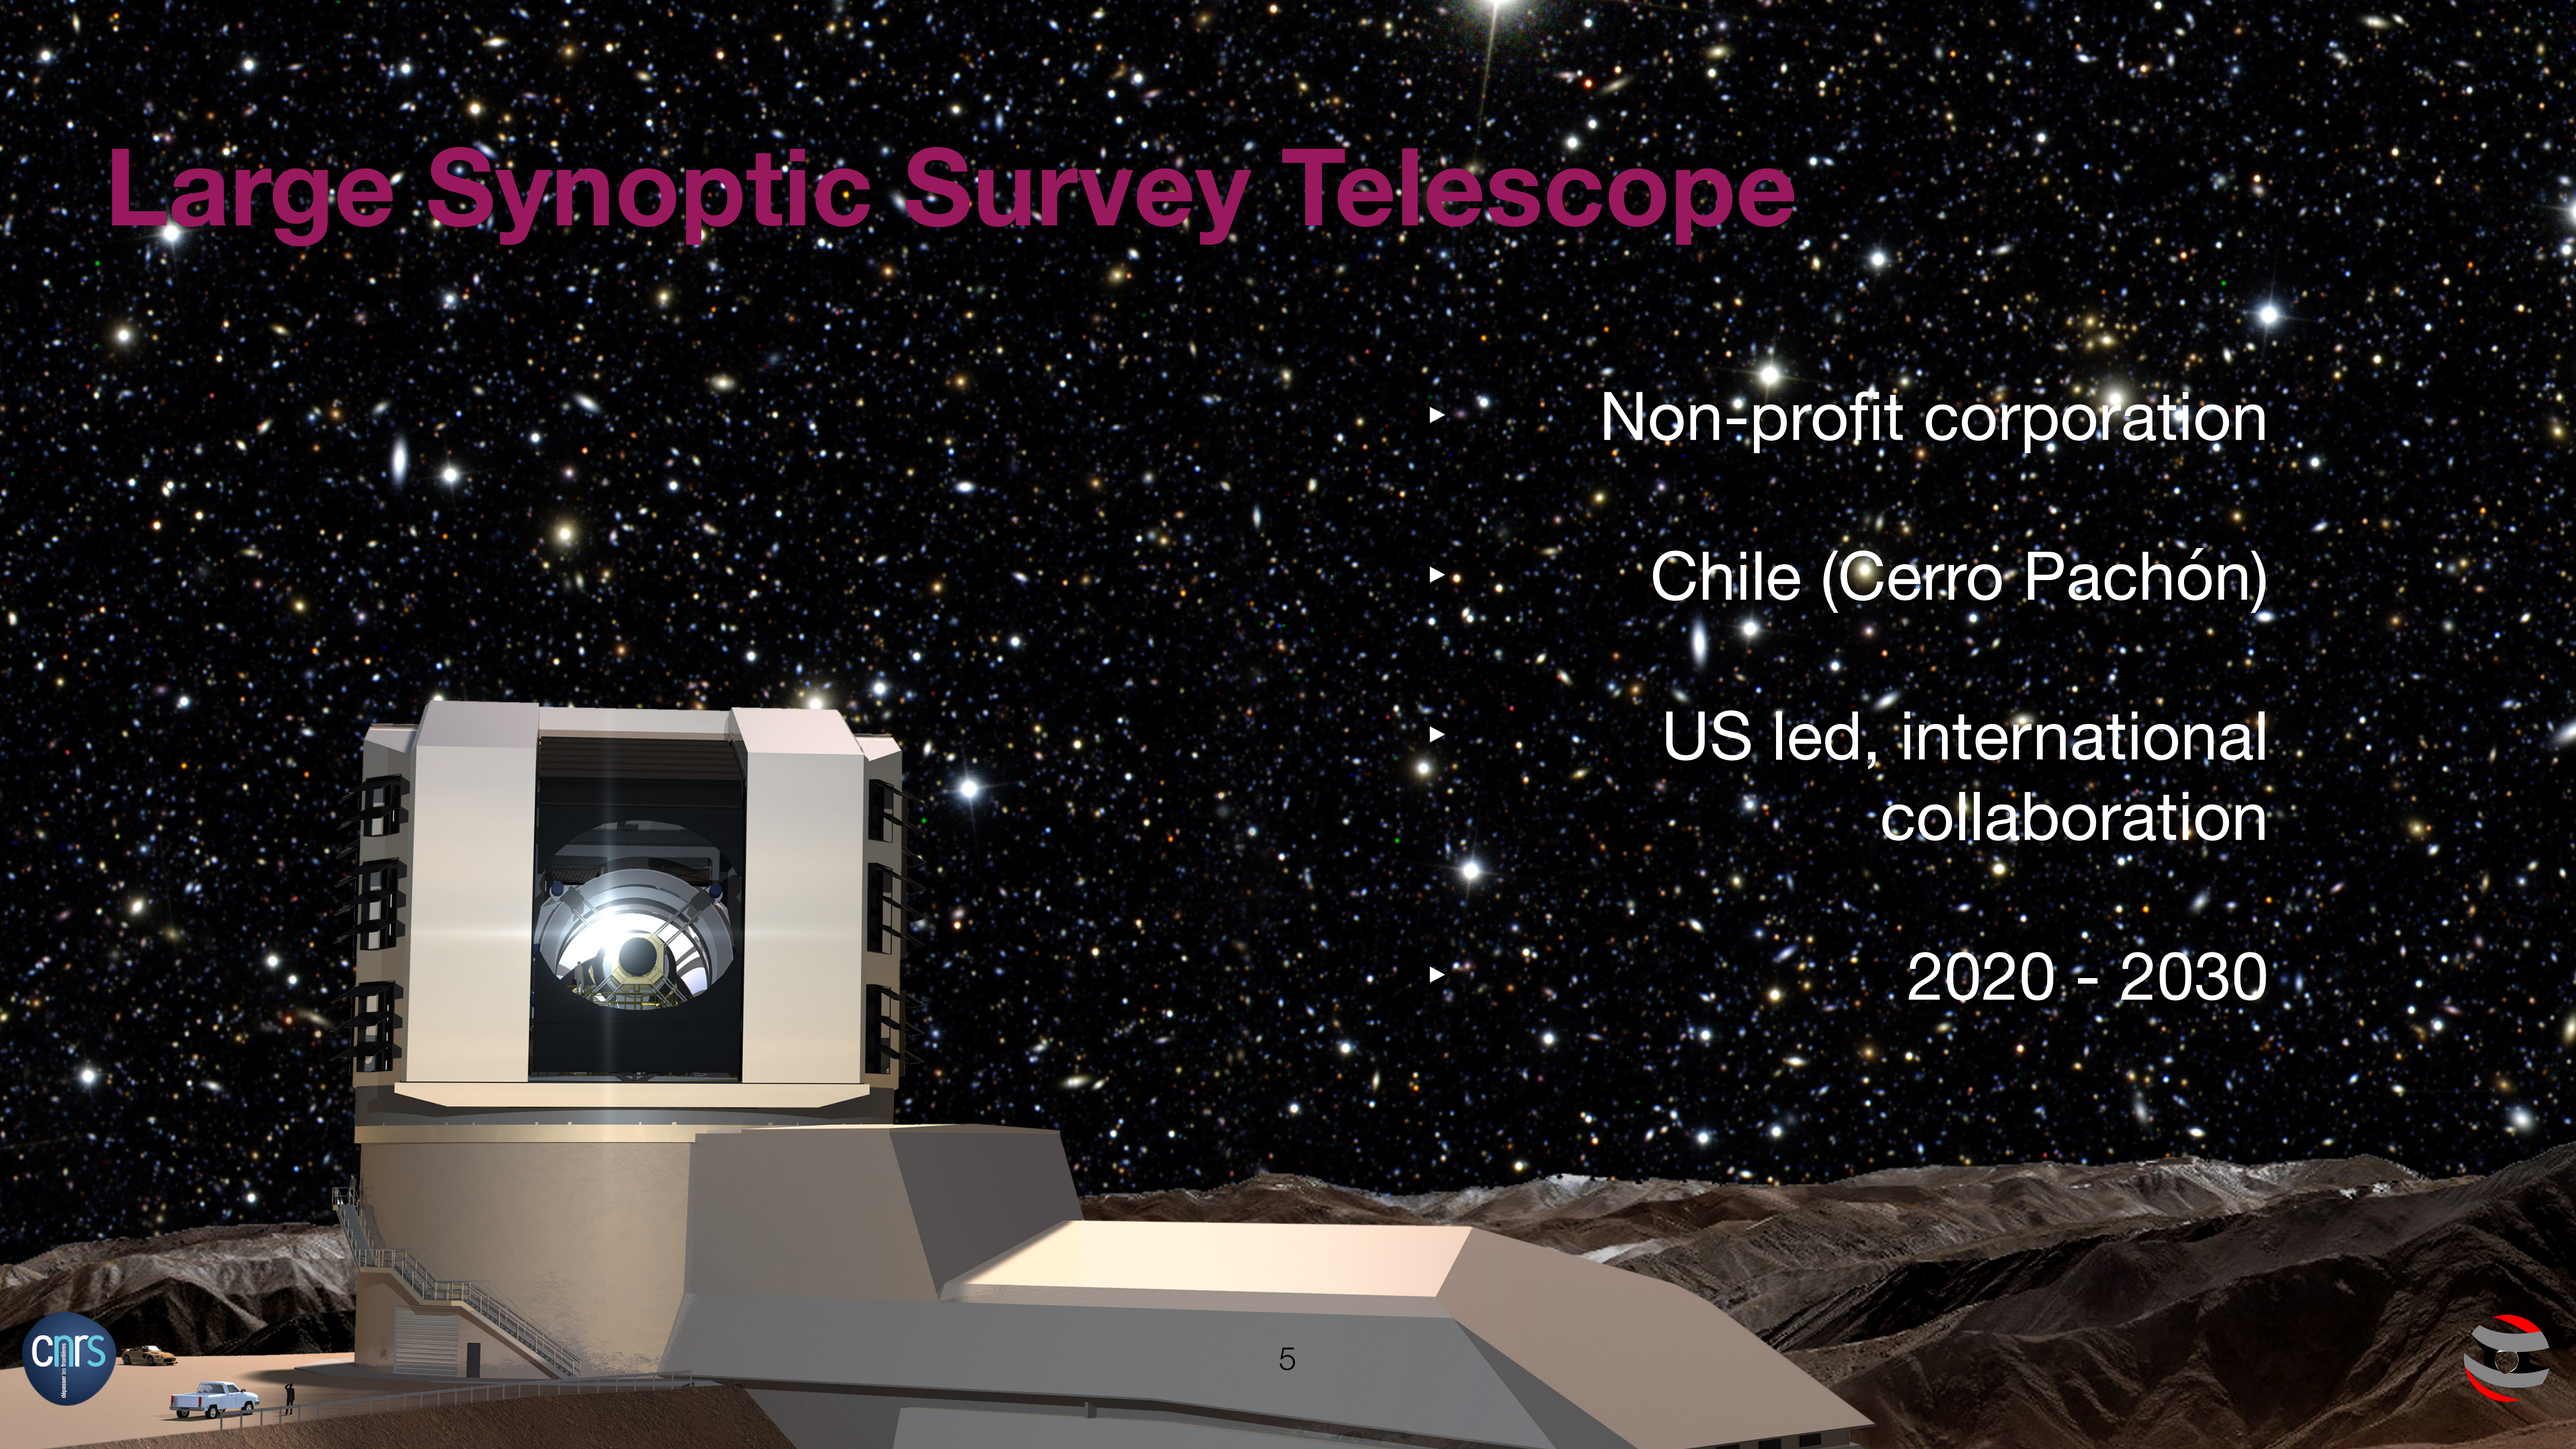

Million pieces puzzle

- 1 exposure / 30 seconds (3.5 x 3.5 sq.deg)
  - Full sky above Chile every 3 nights (~50% full sky).
- 3.2 Gpixels camera
  - 15 TB / night (raw images) + calibration + simulations
- 10 years observation.
  - > 40 billion objects detected.

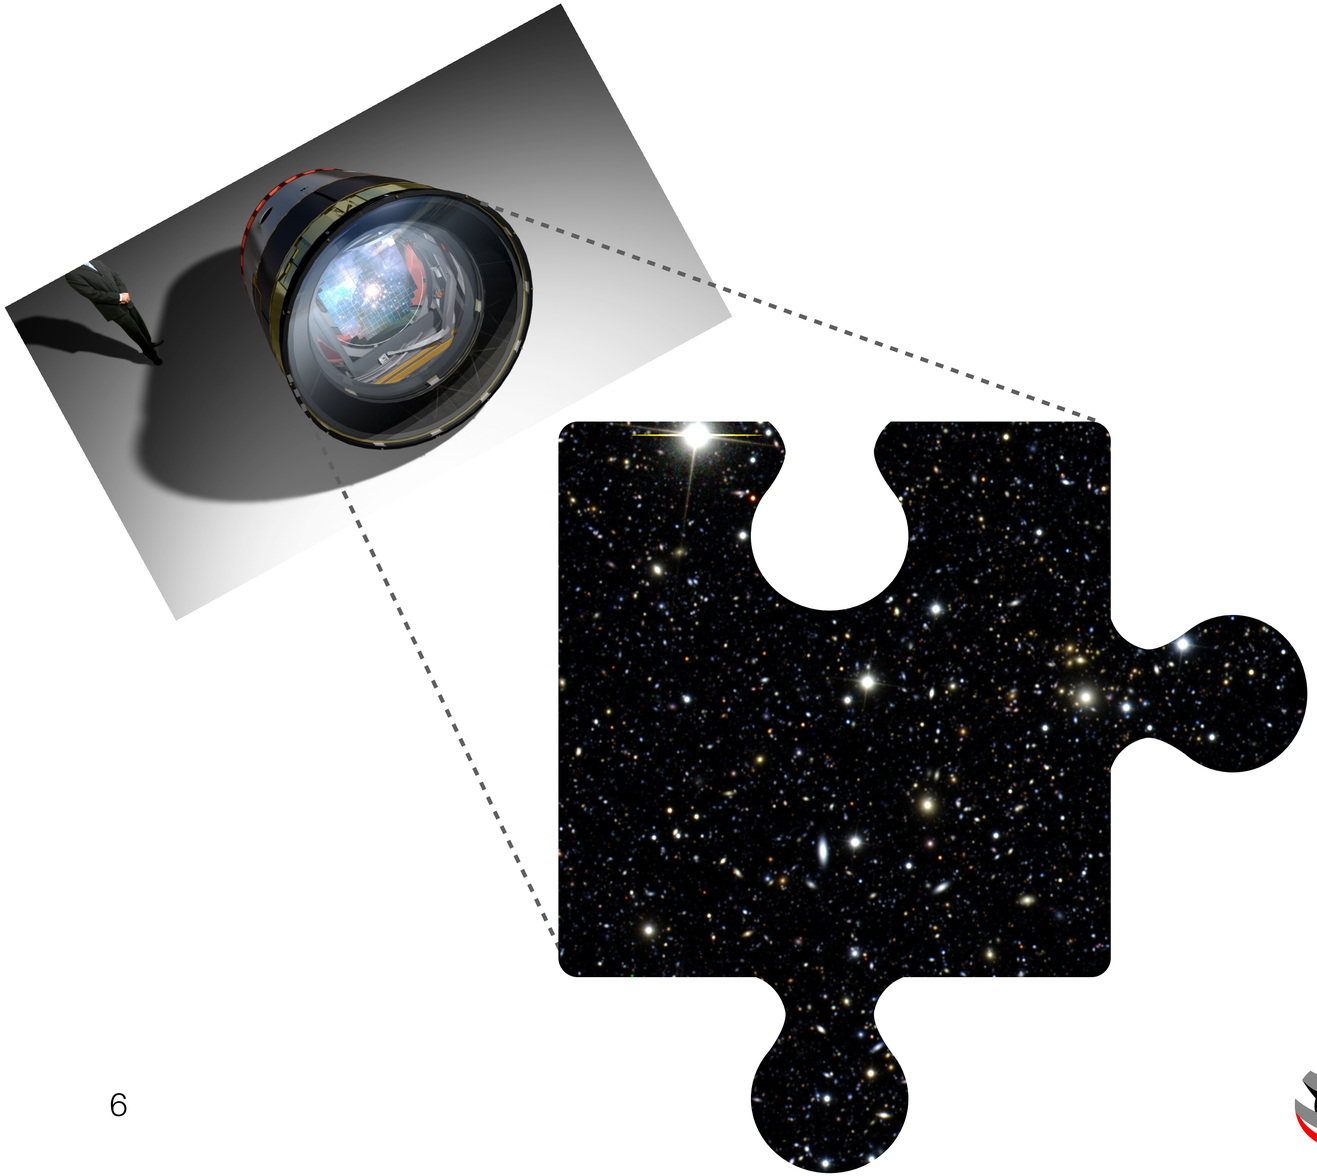

Big data?

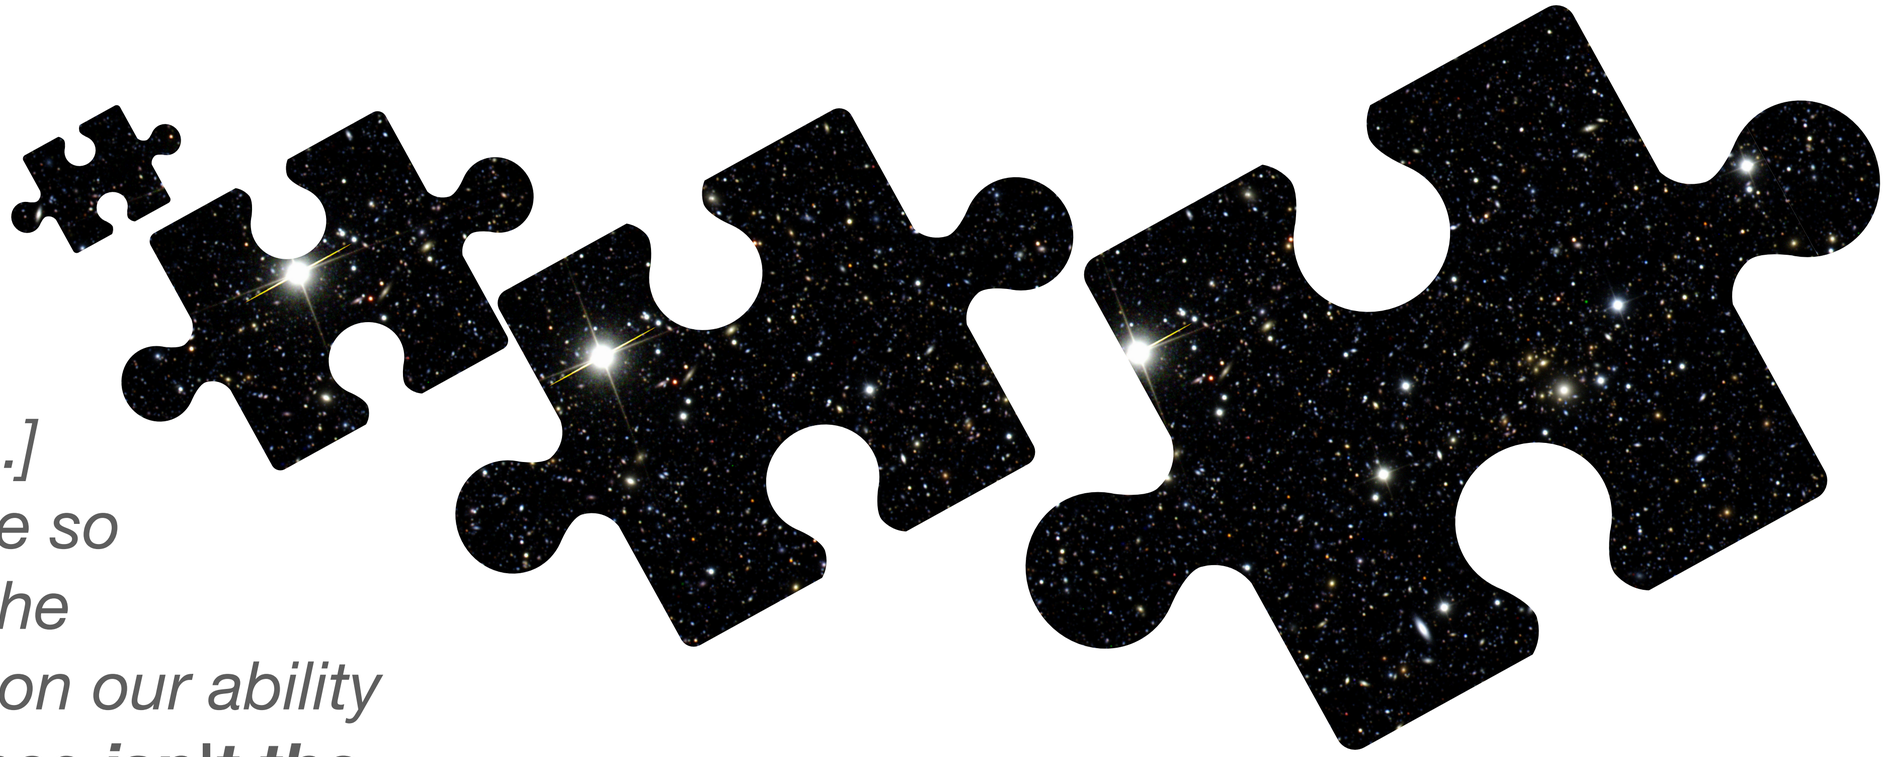

"The data volumes [...] of LSST are so large that the ***limitation*** *on our ability to do science* ***isn't the ability to collect*** *the data,* ***it's the ability to understand*** *[...] the data."* Andrew Connolly (U. Washington)

Astronomical data formats

- Some popular formats in astronomy: **FITS** (1981), HDF (1988).
- Multi-purposes: images, spectra, photon lists, data cubes, or even structured data such as multi-table databases.
- Not designed for serialisation a priori.
- Often use compression.

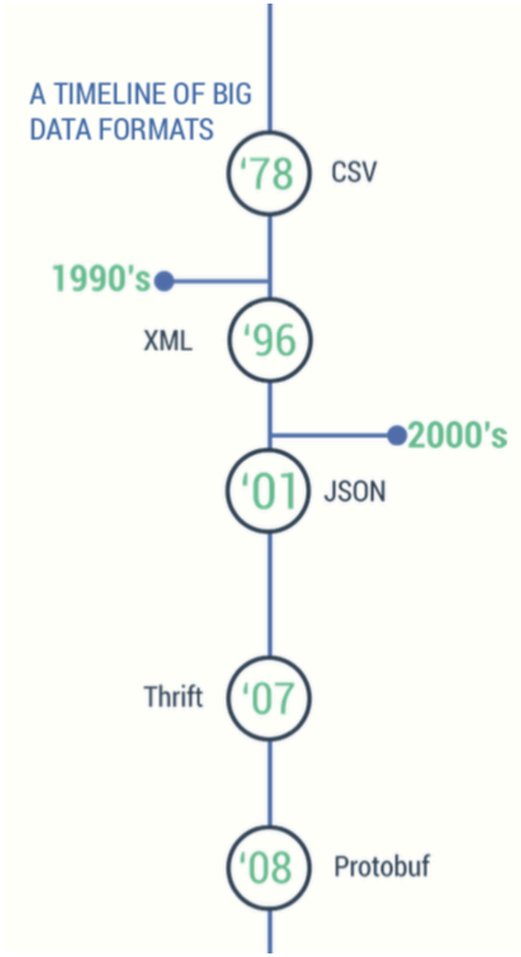

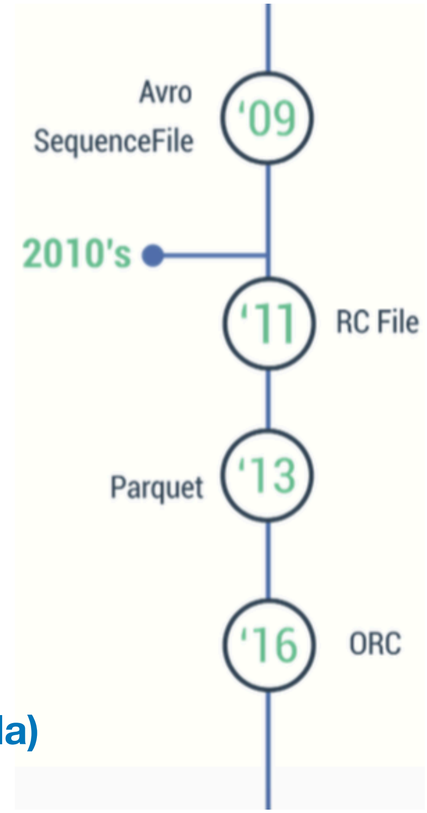

(Nexla)

Flexible Image Transport System (FITS)

- First (latest) release: 1981 (2016).
- Endorsed by NASA and the International Astronomical Union. Maintained by the IAU FITS Working Group.
- Backward compatible
- Set of blocks.1 block: ASCII header+binary data arrays of arbitrary dimension
- Support for C, C++, C#, Fortran, IDL, Java, Julia, MATLAB, Perl, Python, R, and more…

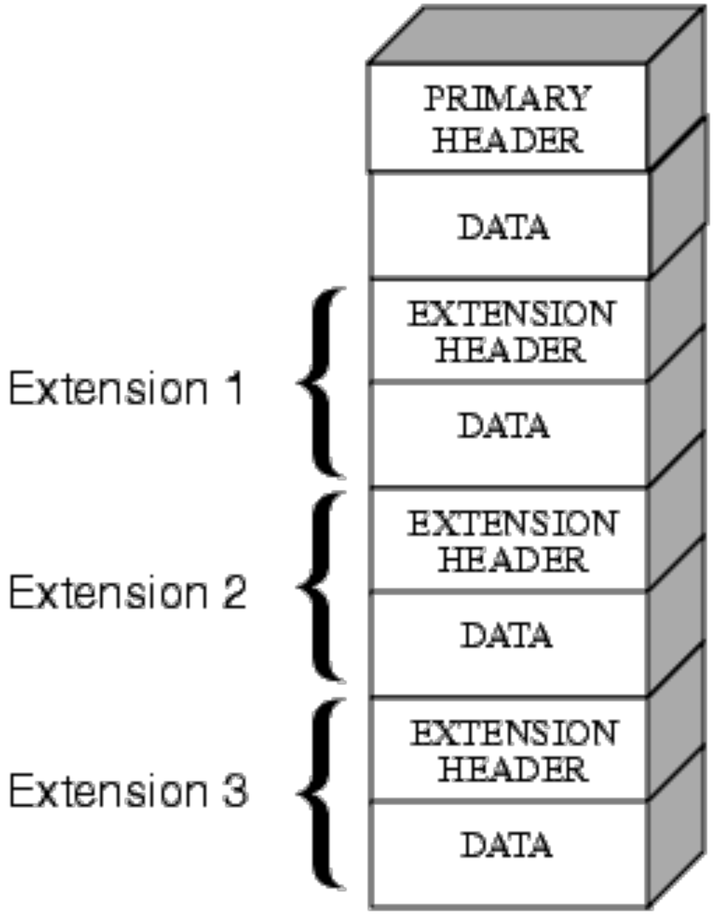

Not only astronomy!

- Vatican Library
- Digital Vatican Library
- FITS and the Vatican Library
- ESA Long-term Data Preservation

ESA AND VATICAN WORK TO PRESERVE HERITAGE DATA

ESA and Vatican Library work together

4 May 2018 Following an agreement signed in 2016, ESA and the Vatican Apostolic Library have presented the preliminary results on how they are using the same method to manage their respective long-term data preservation projects and how this can create new initiatives.

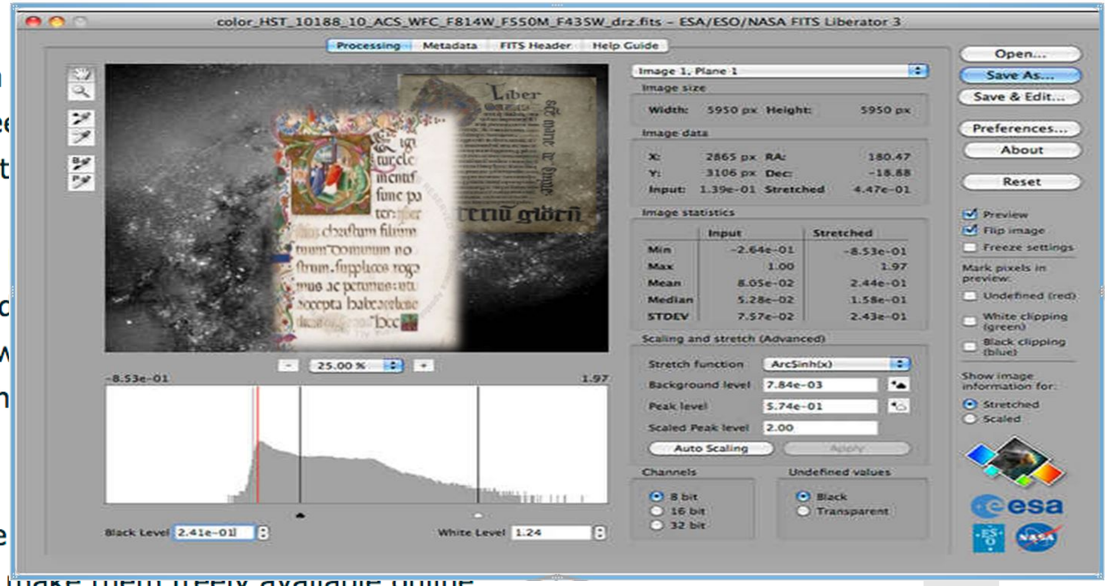

For more than
Library has be
restoring a pat

Now, with the
the Library is c
manuscripts, w
mostly from th
Period.

The aim of the
2010, is to preserve high-resolution images for the long term and make them freely available online through the Digital Vatican Library.

The need for a native Spark connector for FITS

- No Spark connector to read FITS & no FITS support for Scala… :-(
- Previous attempts to use FITS with Spark: read as binary and reconstruct entire block on-the-fly.
  - Partition size is limiting
  - Load balancing can be poor.

Human mind
Spark mind
DFS mind

spark-fits under the hood

https://astrolabsoftware.github.io/spark-fits

- FITS data source for Spark SQL and DataFrames.
- Extend Hadoop `FileInputFormat` class. Images + tables available.
- Schema automatically reads from the FITS header.

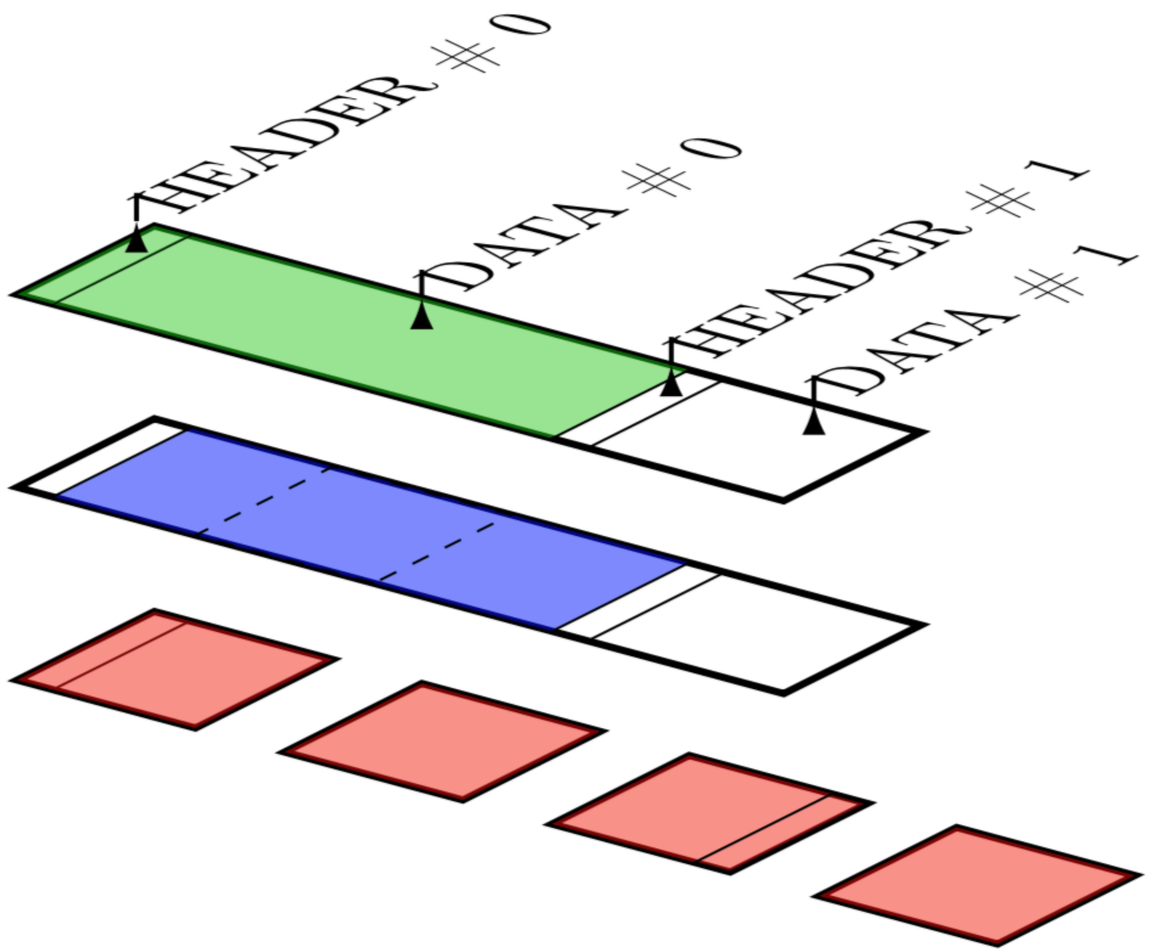

How to use spark-fits

- libraryDependencies += "com.github.astrolabsoftware" % "spark-fits_2.11" % "0.6.0"

```
// Read as a DataFrame a HDU of a table fits.
val df = spark.read
   .format("fits")
   .option("hdu", <Int>)
   .option("columns", <String>)
   .option("recordlength", <Int>)
   .option("verbose", <Boolean>)
   .schema(<StructType>)
   .load(<String>)
```

```
// The DataFrame schema is inferred from
// the FITS header
df.printSchema()
 root
    |-- target: string (nullable = true)
    |-- RA: float (nullable = true)
    |-- Dec: float (nullable = true)
    |-- Index: long (nullable = true)
```

Benchmarks (probably unfair as any bench)

- Spark 2.3.1 / Hadoop 2.8.4
- 1.1 billion rows, 153 cores, 100 iterations each
- No attempt to optimise anything…

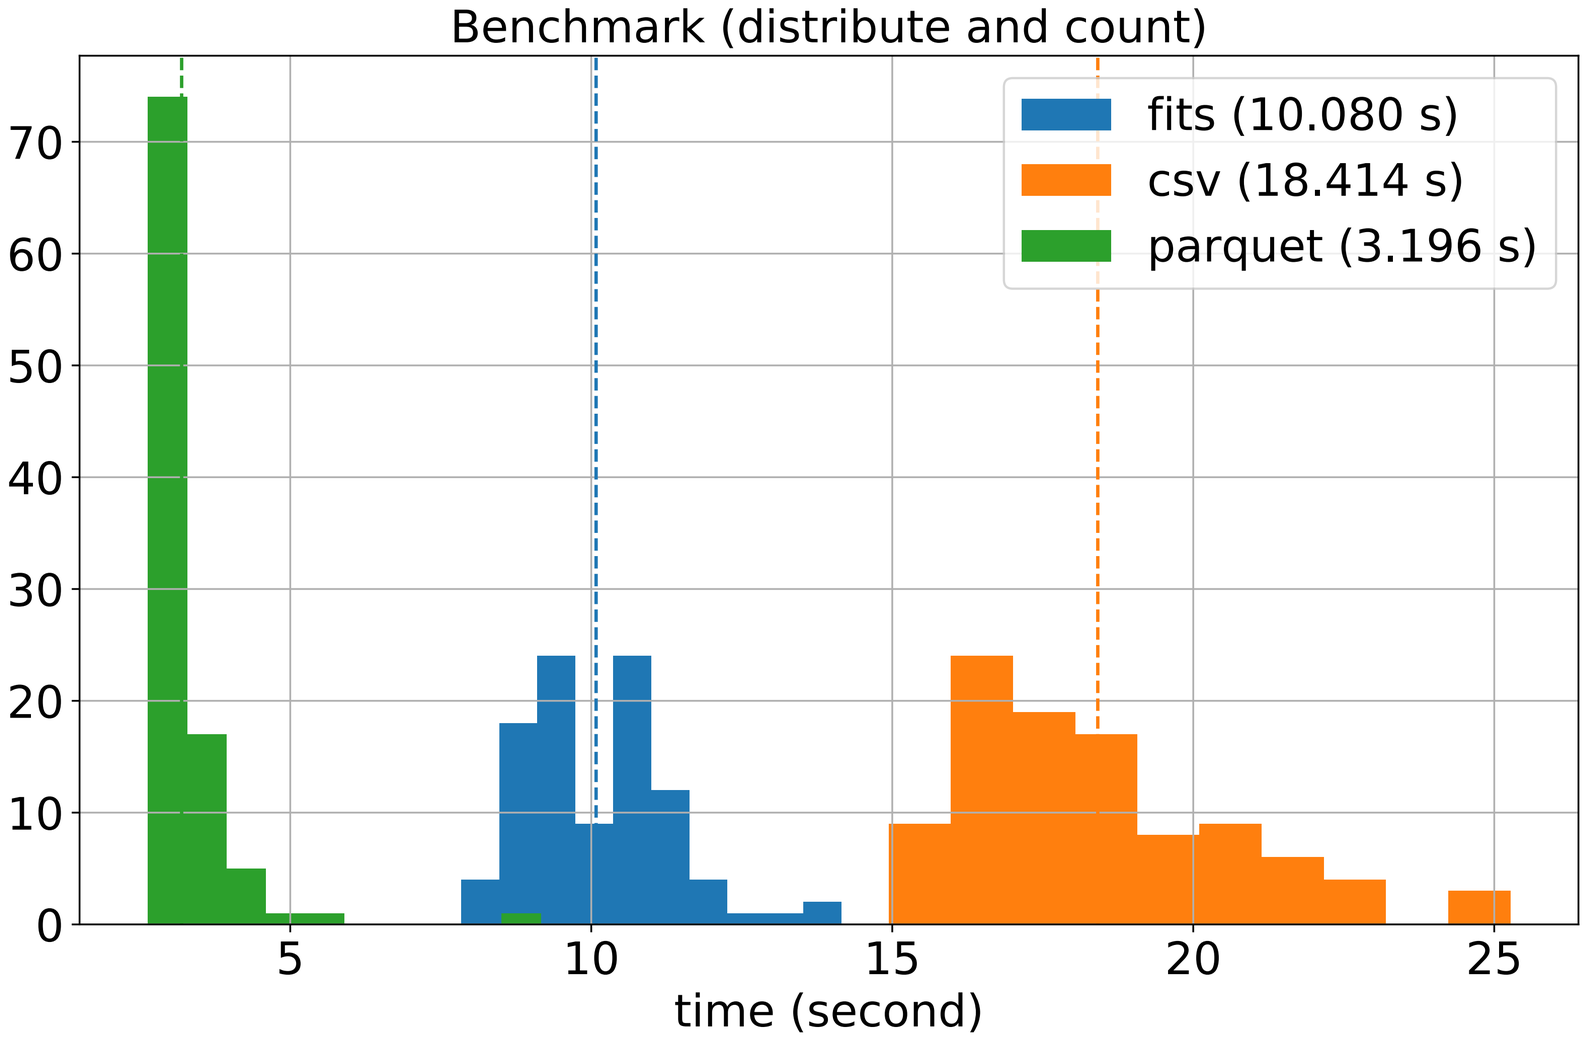

Manipulating large data sets

LAL cluster (153 cores) – up to 6.10^{10} galaxies!

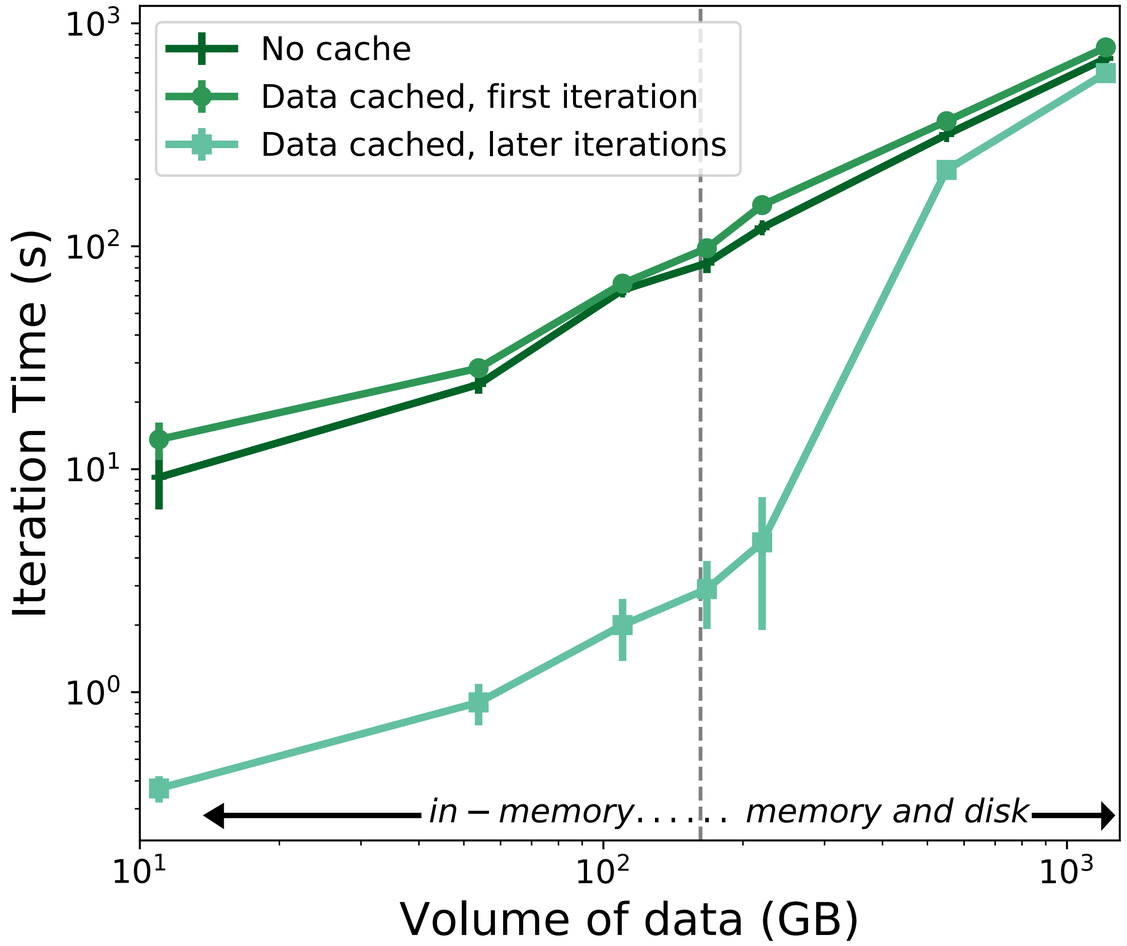

1.2 TB, 1280 cores (@NERSC)

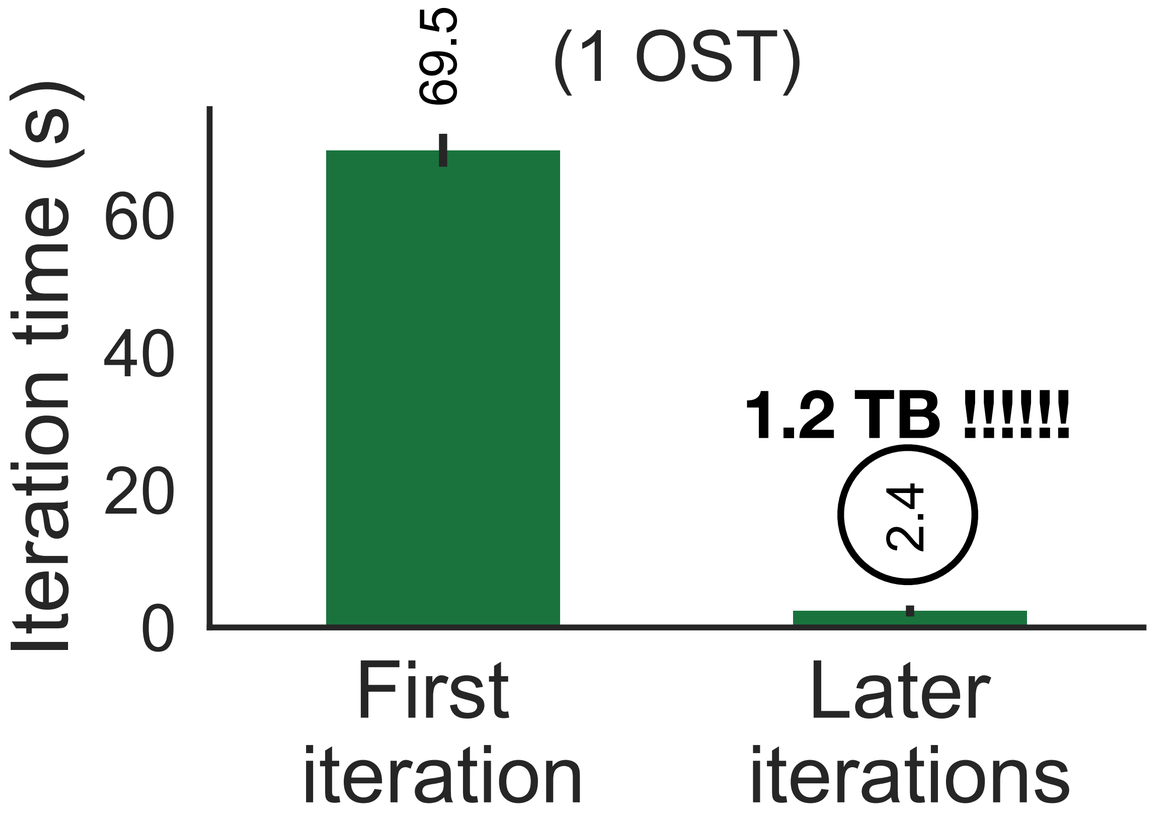

2D vs 3D

- Manipulating 2D data: Geotrellis, Magellan, Geospark, GeoMesa, …
- Very little about 3D!
- Need for e.g. astronomy, particle physics, meteorology.

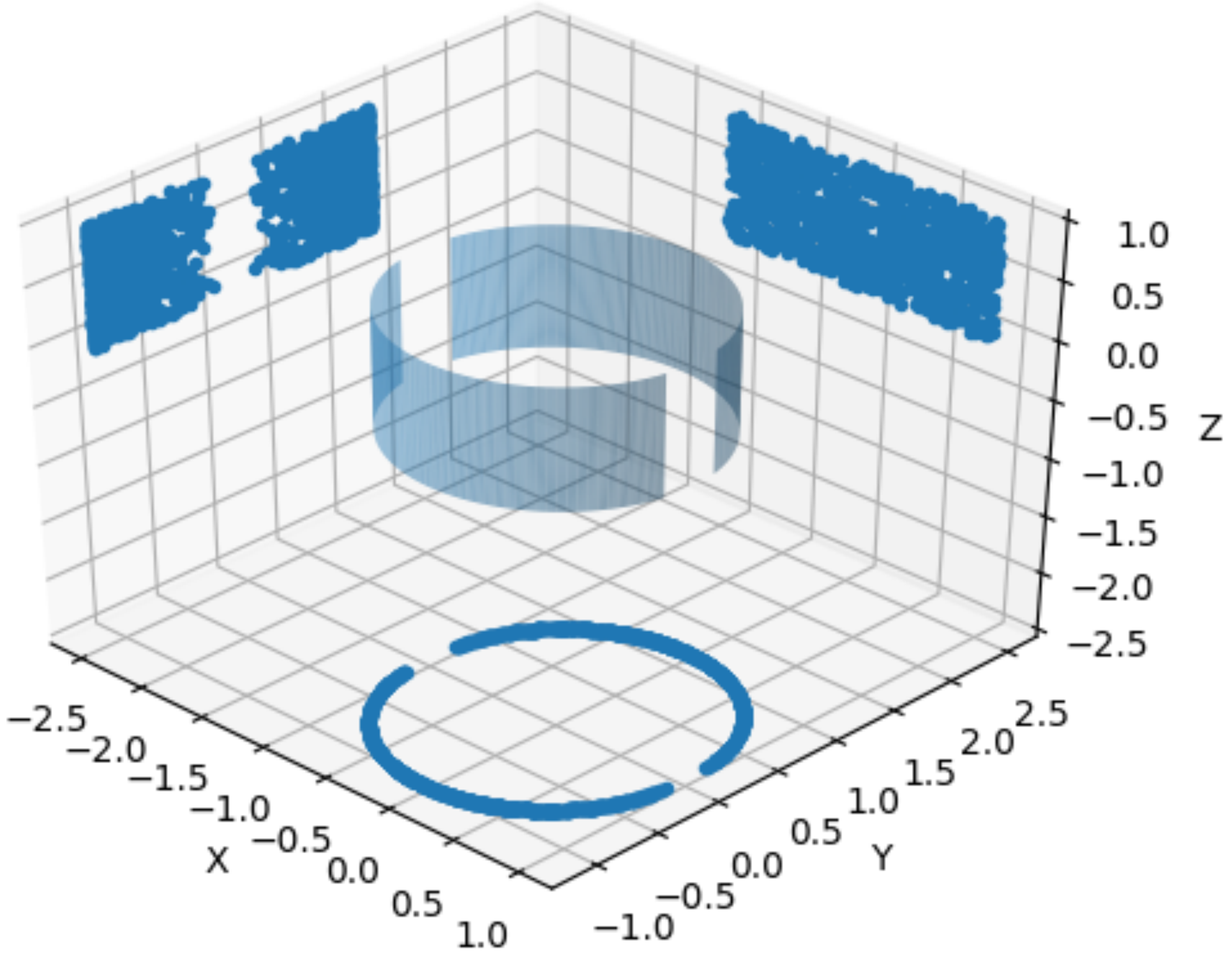

Manipulating 3D spatial data: spark3D

https://astrolabsoftware.github.io/spark3D

- **spark3D**: Starting project.
- GSoC 2018 support.
- 3D partitioning, spatial queries.
- Ongoing: pythonisation (py4j), RDD -> DF, vizualisation, ML and data mining (e.g. DBSCAN)

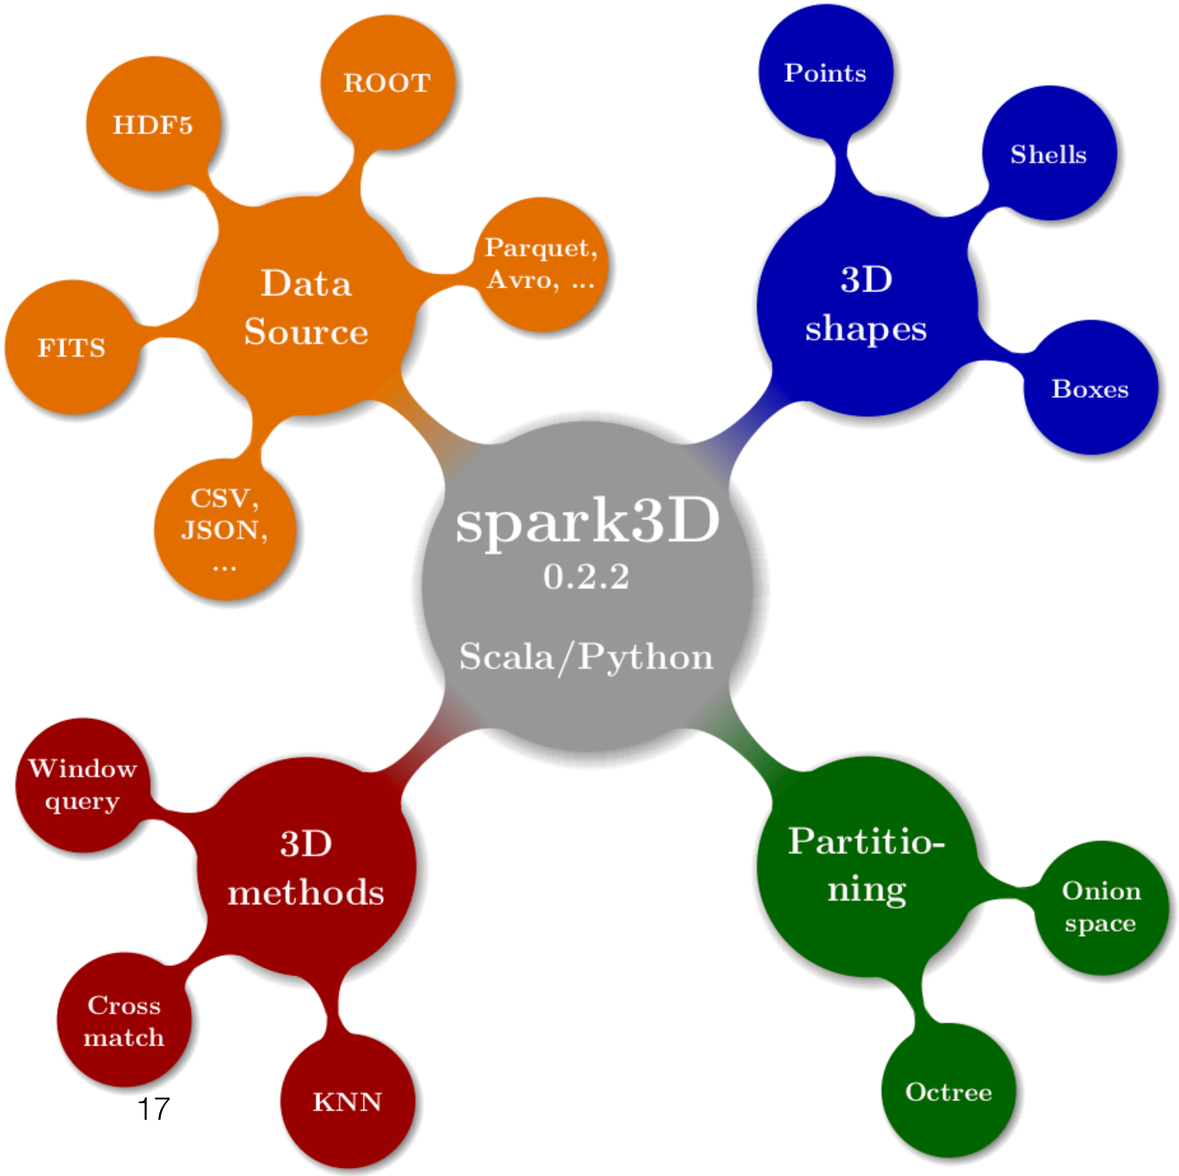

spark3D: re-partitioning

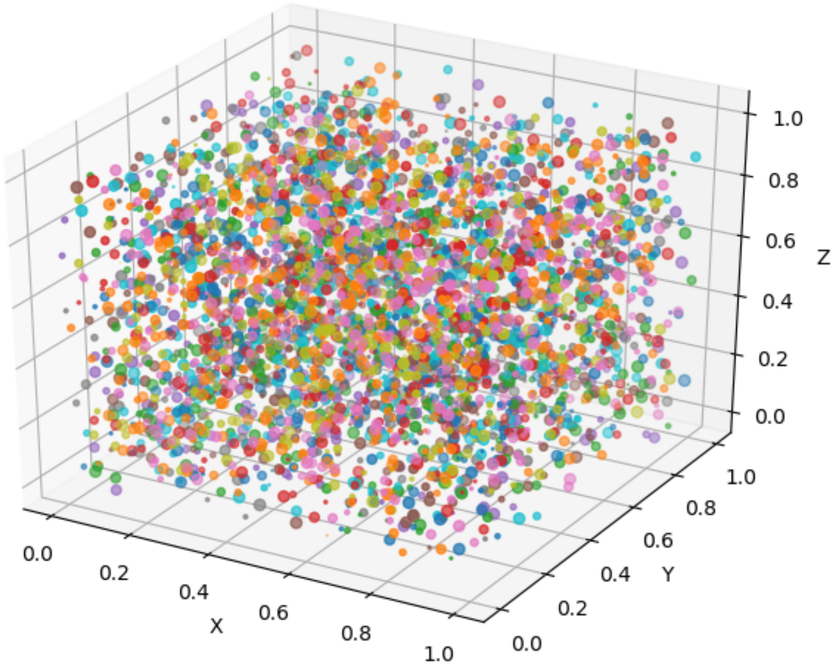

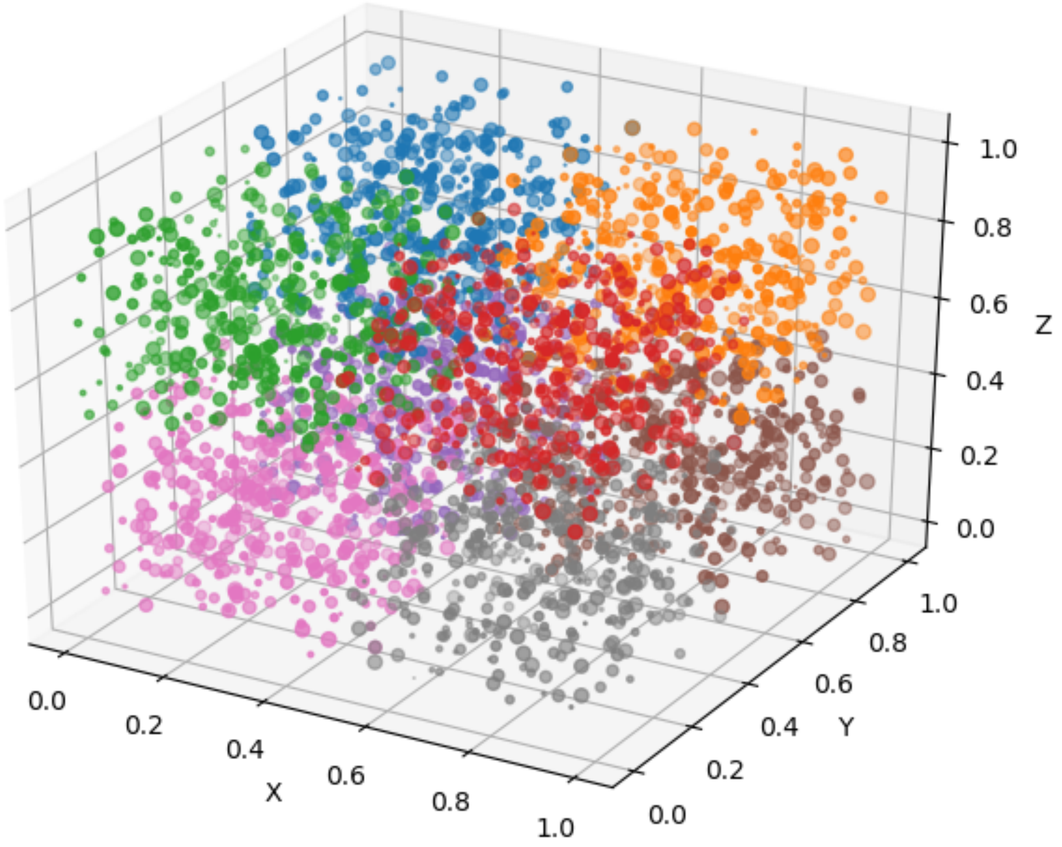

spark3D: K Nearest Neighbours

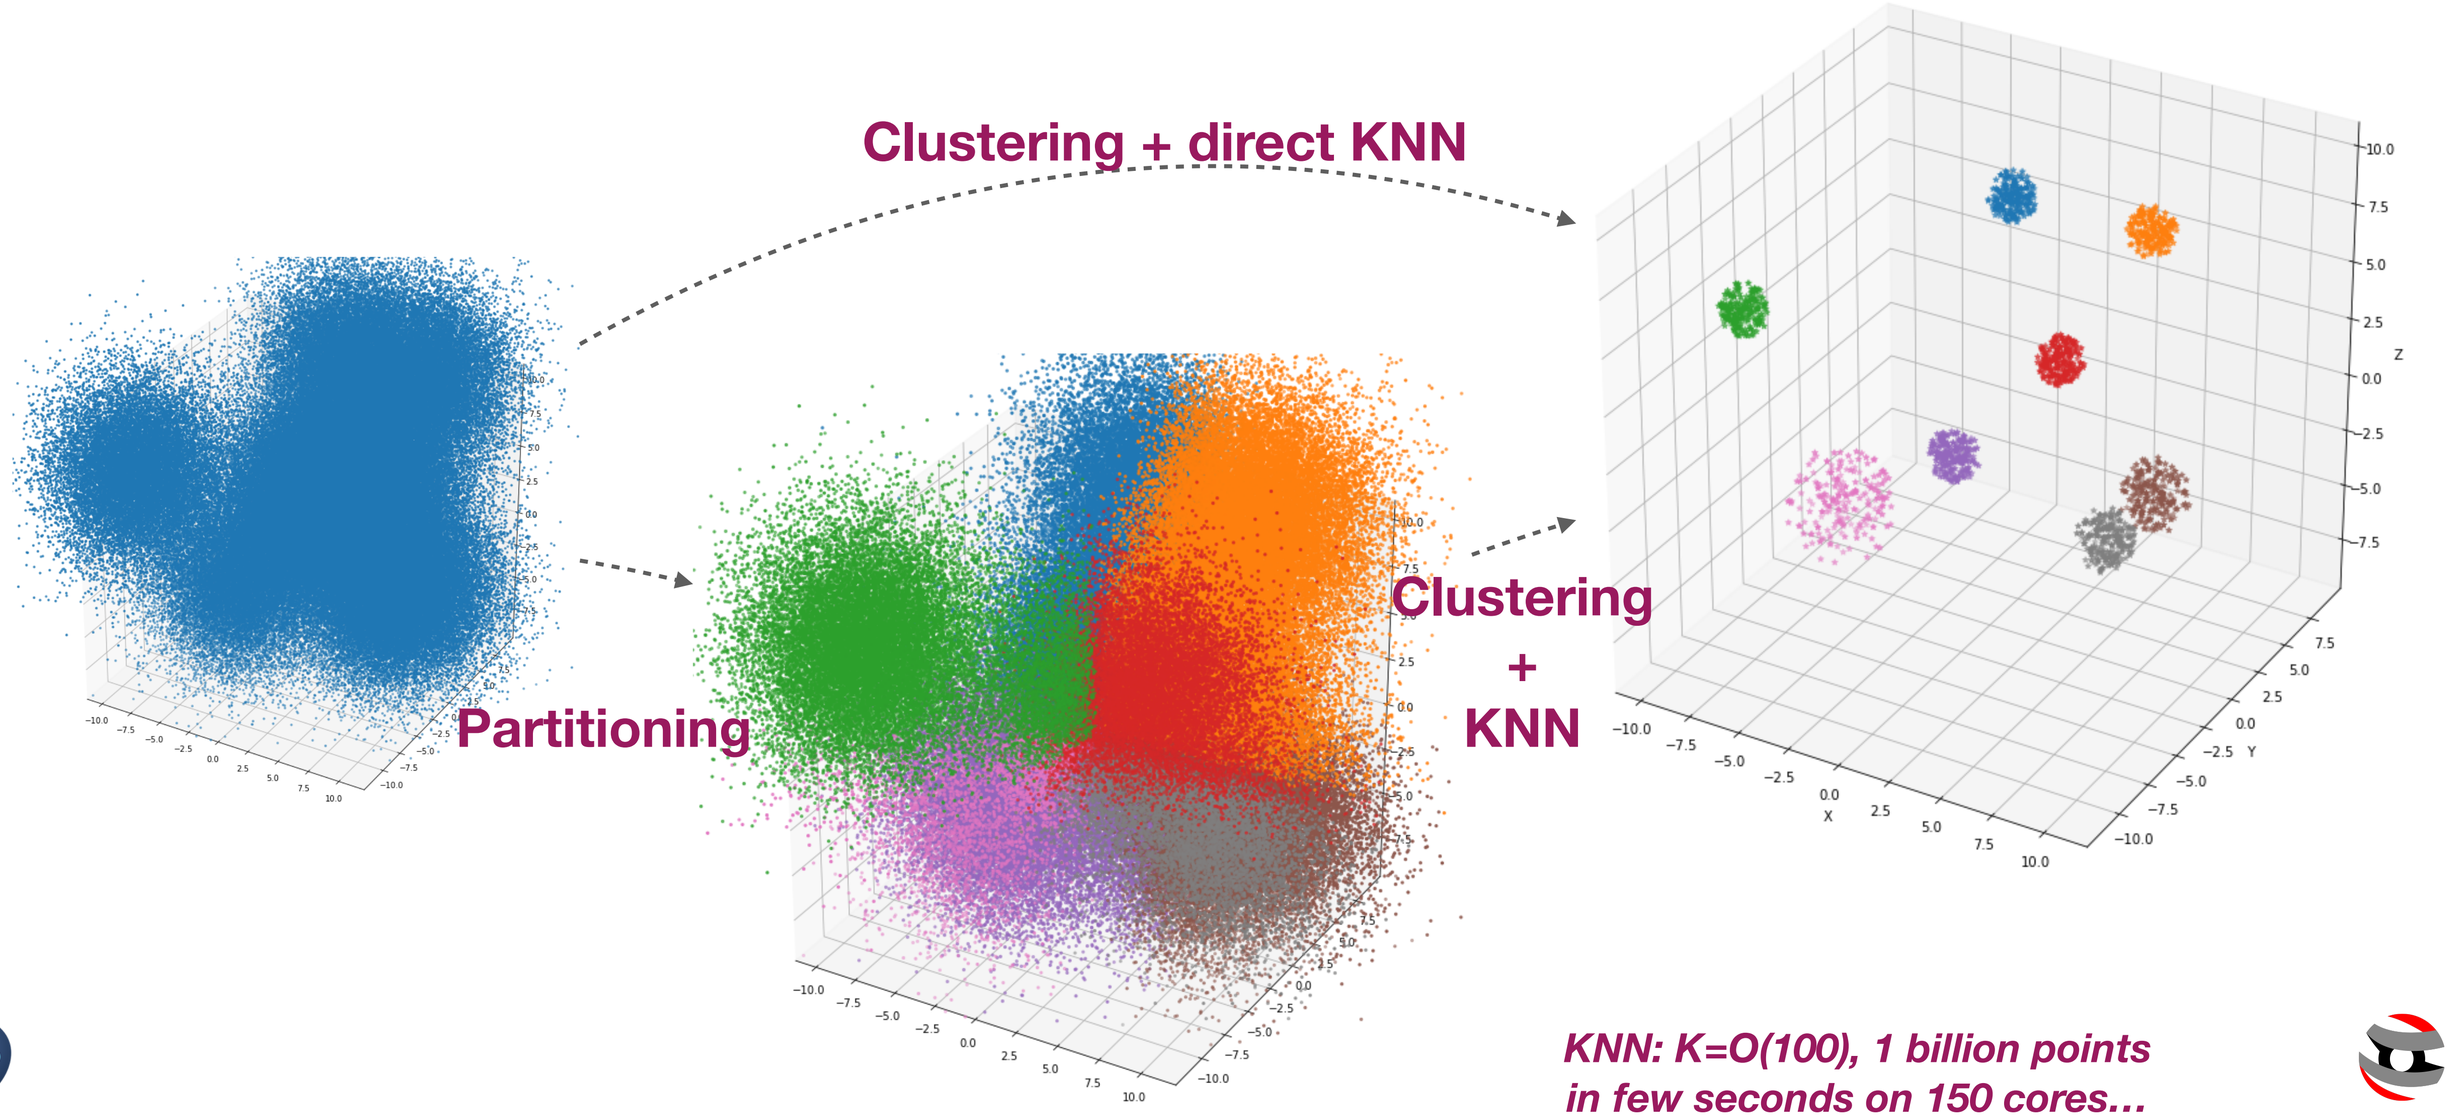

KNN: K=O(100), 1 billion points in few seconds on 150 cores…

When two worlds collide

https://astrolabsoftware.github.io

- Astronomy community in brief*: C++/python oriented, distributed computing is pretty much a theoretical concept: HPC oriented, MPI/multithreading fan.
- Bringing Spark as a *production tool* in this context is a challenge.
- LSST is our incubator.
- AstroLab Software is a project to gather community efforts, and to provide advanced software tools.

*by no means true, certainly cartoonish

At the interface: Apache Spark on HPC

- Little penetration of Spark in HPC domain, but lots of efforts.
- HPC is a serious option for e.g. hybrid workload or heavy shuffle.
- Cori @ NERSC (Cray XC40)
- Apache Spark on Shifter (Docker)

500 million rows, heavy shuffle

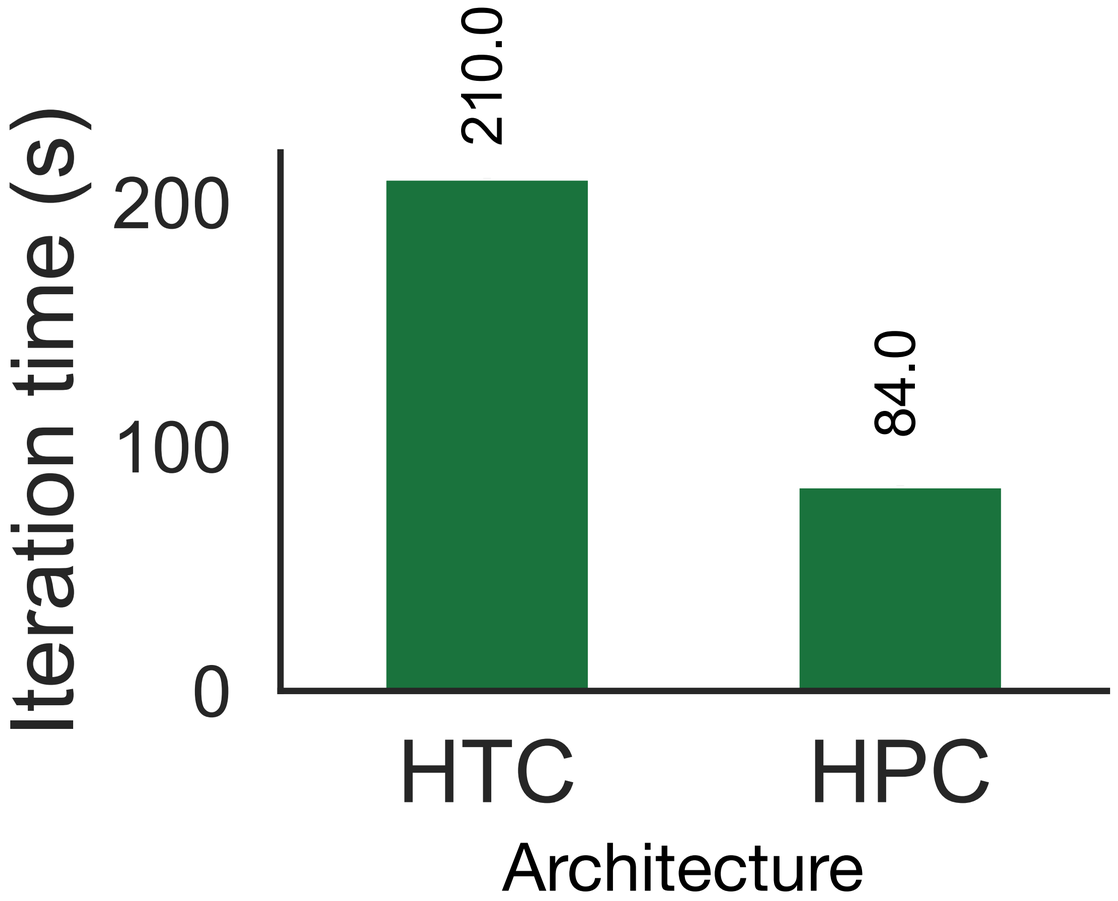

Thanks

- AstroLab Software: https://astrolabsoftware.github.io/
- spark-fits, spark3D, and more!
- You have a big data project in mind? You want to contribute to astronomy? Come talk to us!

Spark @ LAL