Apache Spark, astronomy & LSST DESC

Julien Peloton and Spark @ LAL
CNRS, Laboratoire de l'Accélérateur Linéaire

Christian Arnault
Guy Barrand
Jean-Eric Campagne
Michel Jouvin
Julius Hrivnac
Julien Peloton
Guillaume Philippon
Stéphane Plaszczynski
Adrien Ramparison

Spark@LAL

Apache Spark in one slide

- Cluster computing framework, started in 2009 at UC Berkeley.
- Improve on and extend MapReduce (2004) and Hadoop (2006).
- Open source license (Apache 2.0), mainly written in Scala (Python, R, Java API as well).
- Used by +1000 companies over the world. Little penetration in academia (HEP, biology, astronomy).

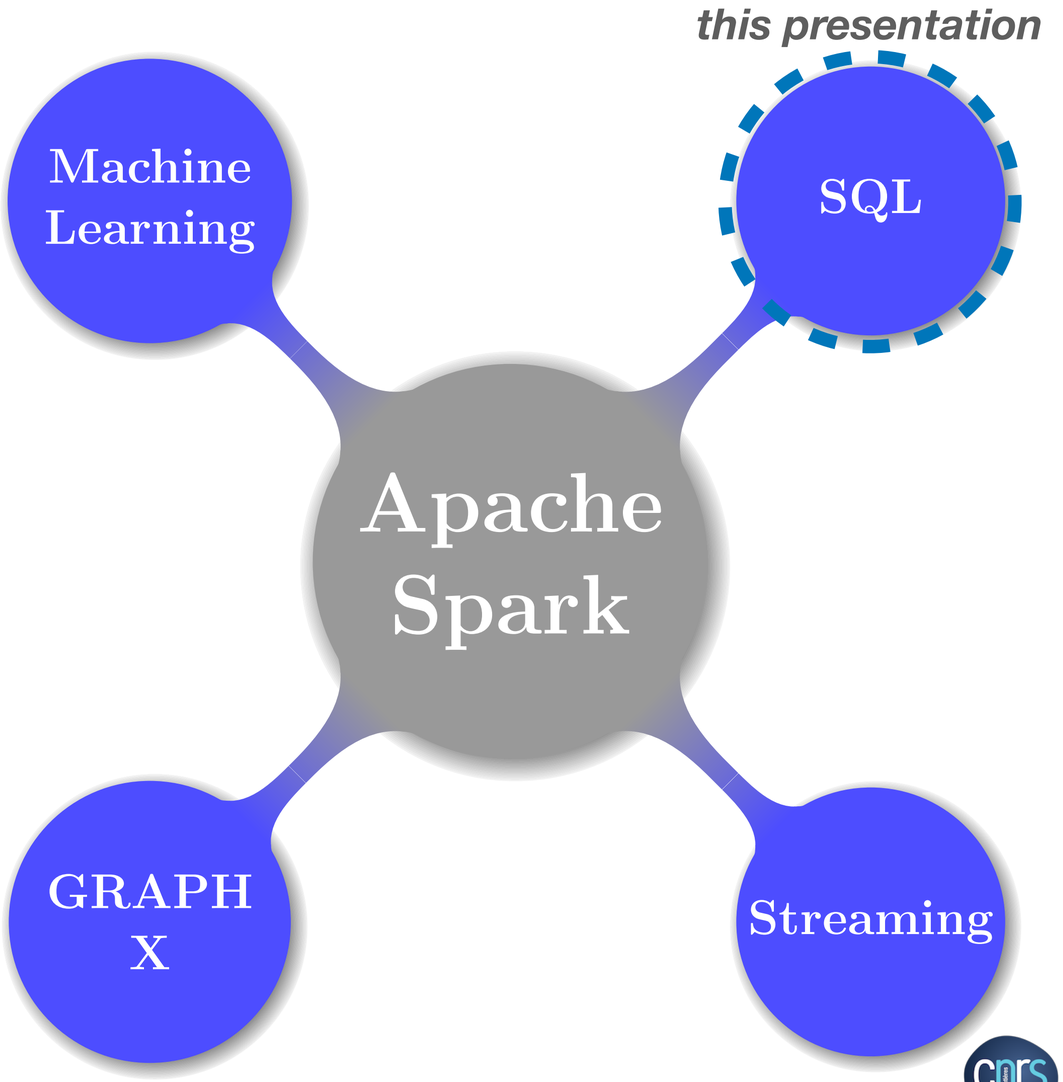

Why Apache Spark for LSST?

- The volume of data collected by LSST every night will be **huge**. The volume of simulation data required will be **huge**. The volume of byproducts will be **huge**.
- But… Big data communities deal with such data volumes (and even more!) for many years.
- Apache Spark proved useful in many contexts and it is an efficient framework to tackle big data problems.
- Apache Spark is designed for : fast prototyping AND interactive data exploration AND production mode.

FITS Data Source for Apache Spark

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

- FITS data source for Spark SQL and DataFrames.
- Image + bintable HDU available. Tested up to 1 TB with linear scaling.
- Computing and Software for Big Science: https://doi.org/10.1007/s41781-018-0014-z (1804.07501)
- Performances (IO throughput) comparable to other built-in Spark connectors.

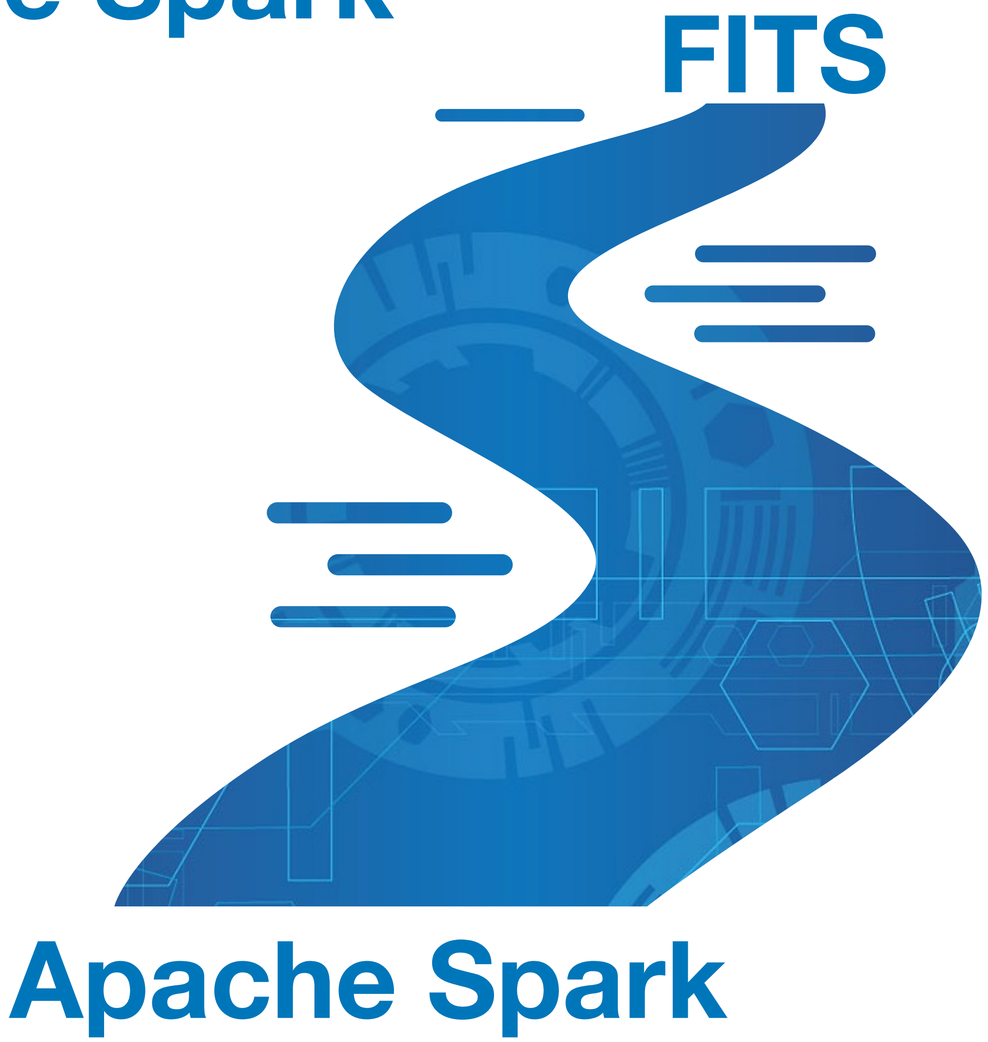

Manipulating 3D spatial data: spark3D

https://astrolabsoftware.github.io/spark3D

- Starting project, under rapid development.
- GSoC + HSF 2018 support.
- 3D distributed partitioning, distributed spatial queries.
- Ongoing projects incl: pythonisation (py4j), vizualisation, machine learning and data mining.

Manipulating 3D spatial data: spark3D

https://astrolabsoftware.github.io/spark3D

- Neighbour search (KNN):
  - 6 billion galaxy positions, K=1000 for a single point **in O(1) sec.**
- Catalog cross-match:
  - 6 billion x 190 million galaxy positions **in O(10) sec.**
- Cluster search (RP-DBSCAN impl. by Song, H. and Lee, J.):
  - 1.2 billion galaxy positions, all cluster centres found **in O(10) min.**

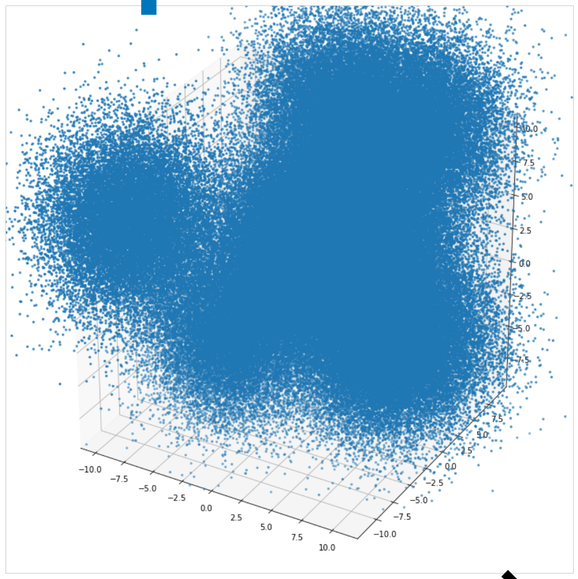

Space partitioning
+
Clustering & KNN

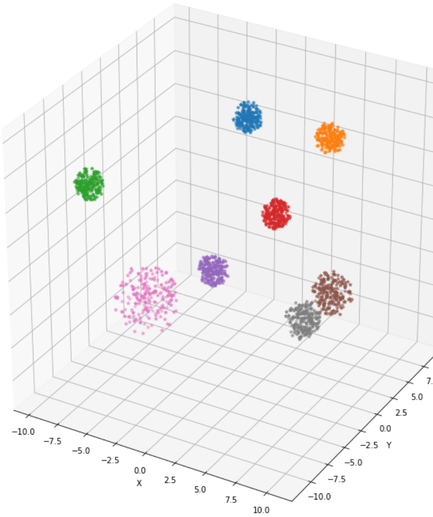

All tests performed on the LAL cluster with 150 cores.

Interfacing Apache Spark with external libraries

- Spark is written in Scala, with also Python, Java and R API.
- External libraries or softwares: rewrite or interface.
- Ongoing projects on: sextractor (E. Bertin), Spectractor (J. Neveu), visualisation tools (G. Barrand with inexlib_py).

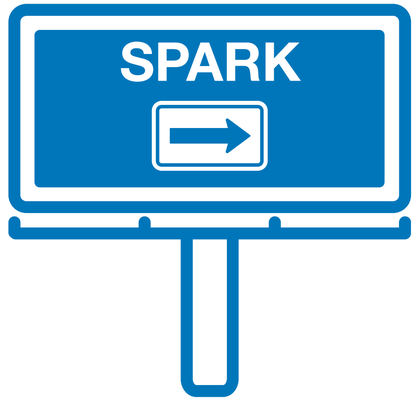

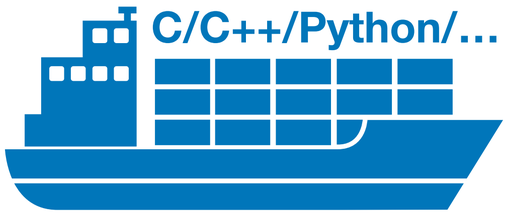

AstroLab Software

- AstroLab Software
  - Project started in 2018 to gather community efforts, and to provide advanced software tools.
  - Big data in astronomy. Focus on Apache Spark.
  - Open source: https://github.com/astrolabsoftware

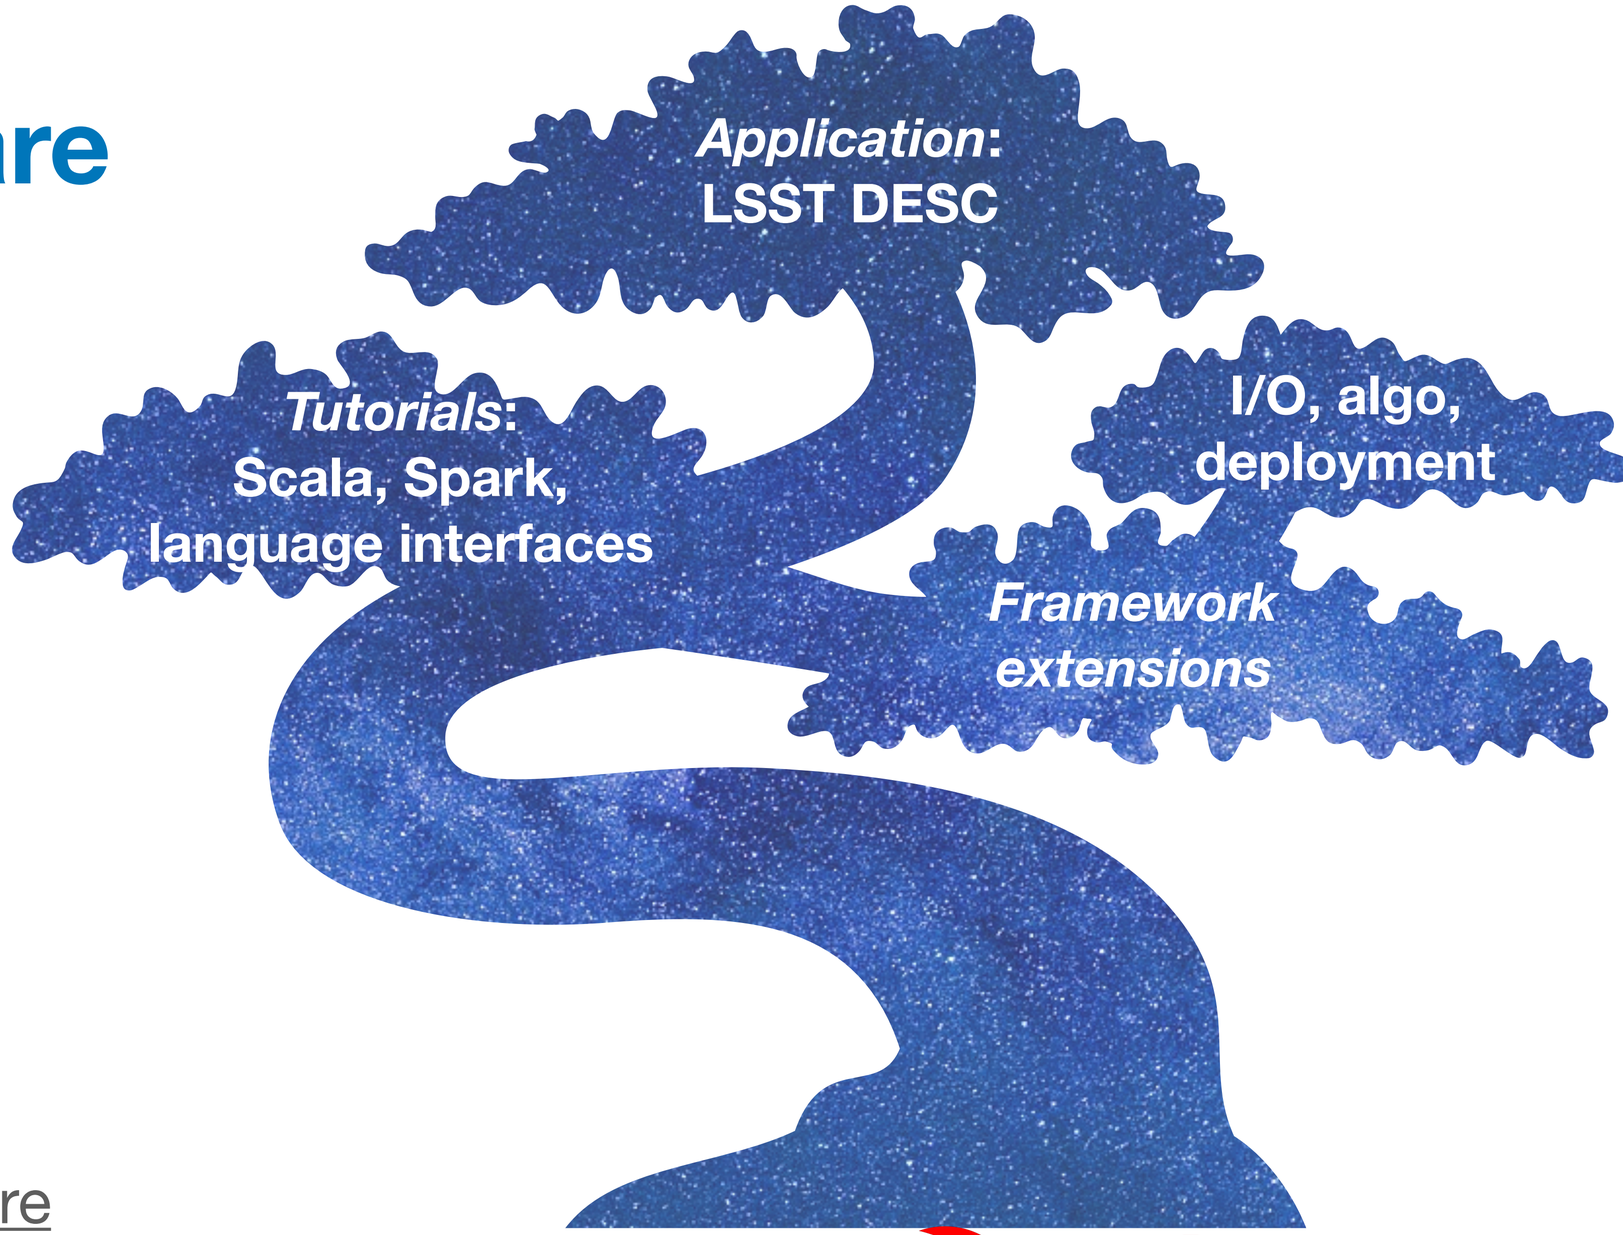

DESC Data Challenge 2 (DC2)

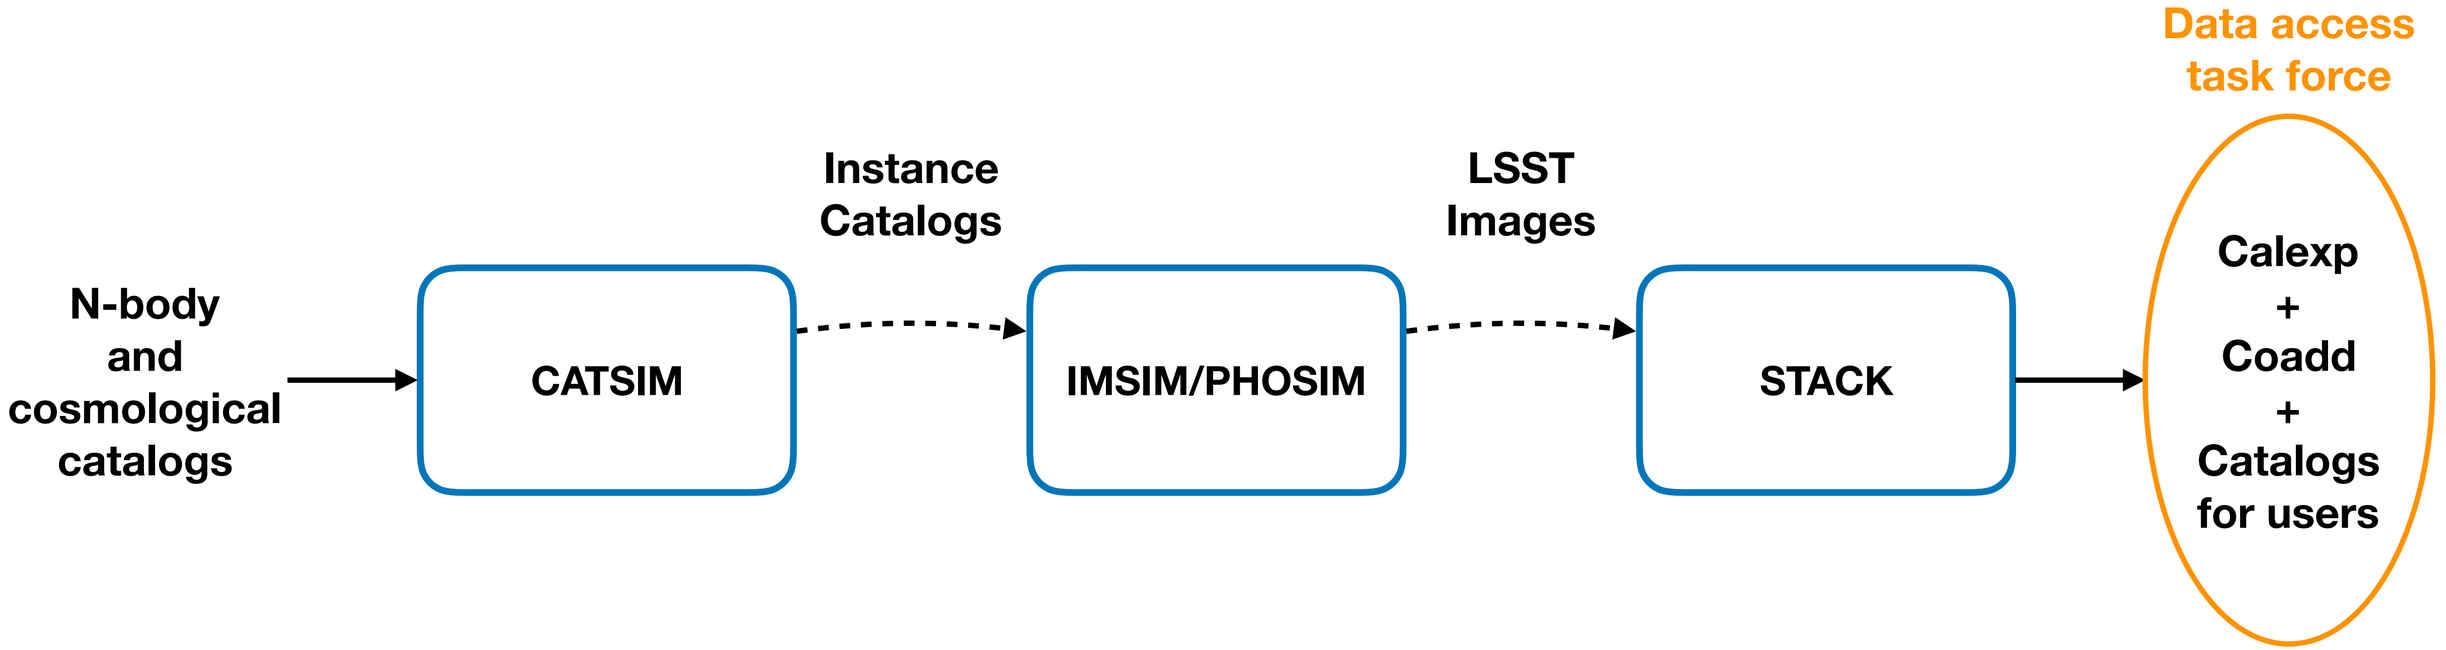

Thanks Dominique, Réza, & Stéphane!

Apache Spark in DESC

- Dedicated DESC repository: https://github.com/LSSTDESC/desc-spark
  - How to set up a DESC python environment with Apache Spark at NERSC (batch/interactive + JupyterLab). Official Kernel soon distributed with the DESC environment.
  - Tutorials to discover Apache Spark.
  - Links to Apache Spark developments in DESC (DC2 data access, 3x2pt, …).
- Frequent discussions with the NERSC consulting group.

Apache Spark in DC2 (Run 1.1p)

Explore DC2

- https://github.com/LSSTDESC/DC2-analysis/pull/27
- Learning objectives:
  - **Load** and efficiently access a DC2 object **catalogs** (DPDD) with Apache Spark
  - **Understand** and have references for the catalog **schema**
  - **Apply cuts** to the catalog using Spark SQL functionalities
  - Have an **example** of quality cuts and simple star/galaxy separation cut
  - **Distribute the computation** and the routine to plot to be faster!

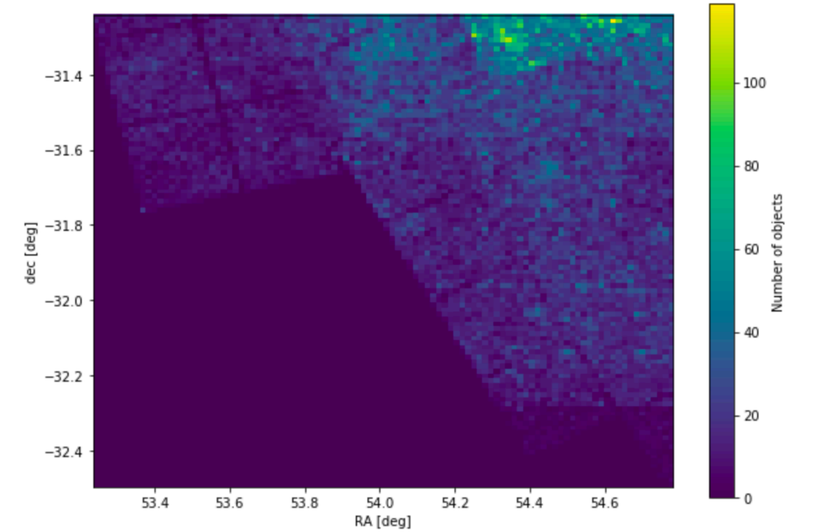

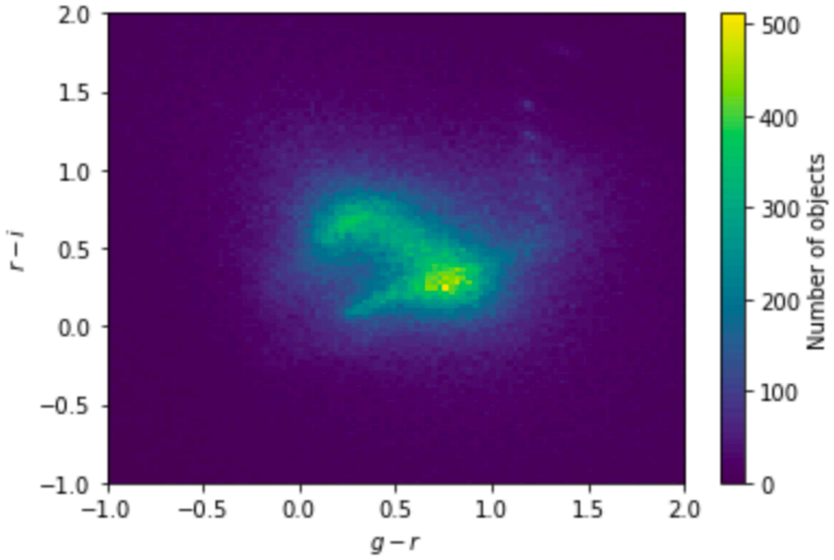

Apache Spark in DC2 (Run 1.1p)

- https://github.com/LSSTDESC/DC2-production/pull/288

- Learning objectives:

  - Introduce Apache Spark and test performance of data manipulations of the static coadd catalogs (DPDD)

  - Detailed benchmarks on Parquet and FITS format.

Data set	#Rows (size GB)	Load time
Parquet (OT)	719,228 (0.43)	393 ms ± 86.2 ms
FITS (OT)	719,228 (0.57)	312 ms ± 59.2 ms
Parquet (AT, Hive)	6,892,380 (4.5)	215 ms ± 102 ms
Parquet (AT, Simple)	6,892,380 (3.6)	181 ms ± 74.1 ms
FITS (AT)	6,892,380 (5.4)	2.78 s ± 1.37 s

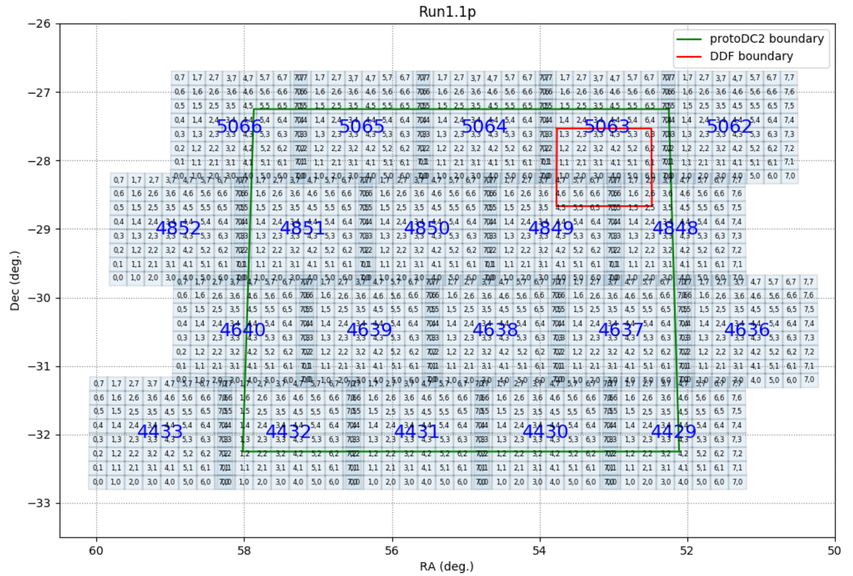

Apache Spark in DC2

- Images and catalogs from image simulation:
  - Run 1.1p: DM outputs + DPDD-like static object catalog
  - Run 1.2p & 1.2i: DM outputs
  - Run 2.0p: DM outputs (just started)
- Production static object catalogs (FITS/Parquet) from DRP
  - https://github.com/LSSTDESC/DC2-production/commits/issues/273
- Participation to 3 x 2 pt

Future (in LSST DESC)

- DC2, more and more! (HackurDC2?)
- Explore end-to-end pipeline(s) for image/catalog processing.
- Bootcamp Apache Spark for DESC members.
- Connect to visualisation tools.
- Streaming capability.
- Tests on Google Cloud?
- …