Institute of Cosmos Sciences, University of Barcelona
29-May to 2-June 2023

6th Barcelona TechnoWeek

Scientific Computing in the Cloud

Cloud use case 2: Astrophysics

Galactic RainCloudS

(Galactic Research in Cloud Services)

Jordi Portell, Óscar Jiménez, Marcel Bernet,
Sergi Bartolomé, Albert Masip, João A.S. Amarante, et al.
Universitat de Barcelona, Institut de Ciències del Cosmos (ICCUB)
Institut d'Estudis Espacials de Catalunya (IEEC)

Universitat de Barcelona
IEEC Institut d'Estudis Espacials de Catalunya
DAPCOM Data services
OCRE Open Clouds for Research Environments
esa gaia
ICCUB Institut de Ciències del Cosmos Universitat de Barcelona
Excelencia María de Maeztu 2020-2023
Unió Europea Fons Europeu Next Generation
Gobierno de España
Plan de Recuperación, Transformación y Resiliencia
Next Generation Catalunya
Generalitat de Catalunya Departament de Recerca i Universitats

The Gaia mission

Global astrometry from Space (2013-2025)

Catalogue of ~2 billion stars (plus other galaxies, Solar system objects, …)

- Positions, distances, motions, brightness, colours, astrophysical params, …
- Worldwide astronomy revolution

Large amounts of data

- Billions of interrelated records
- Processing and exploitation challenge

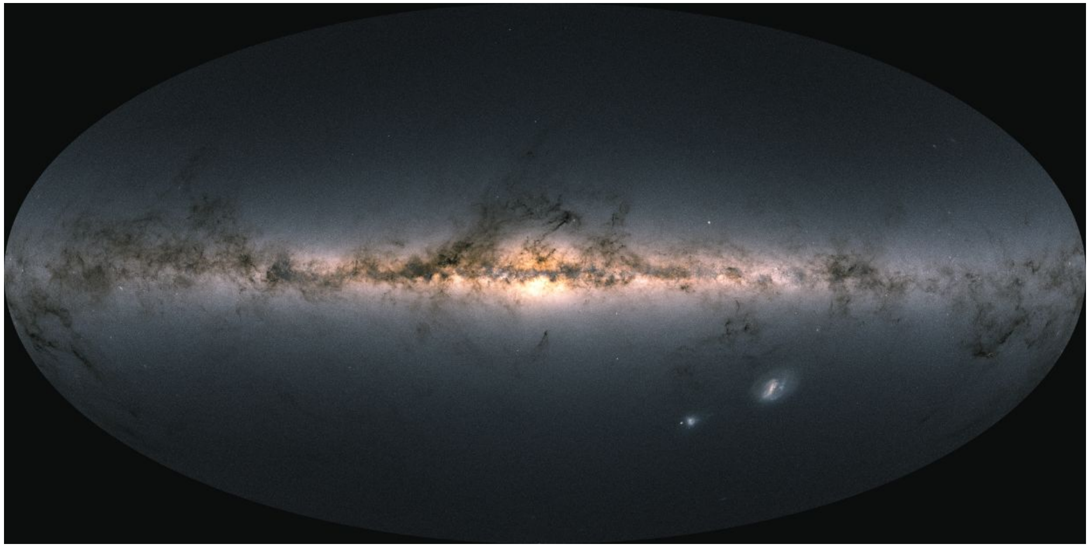

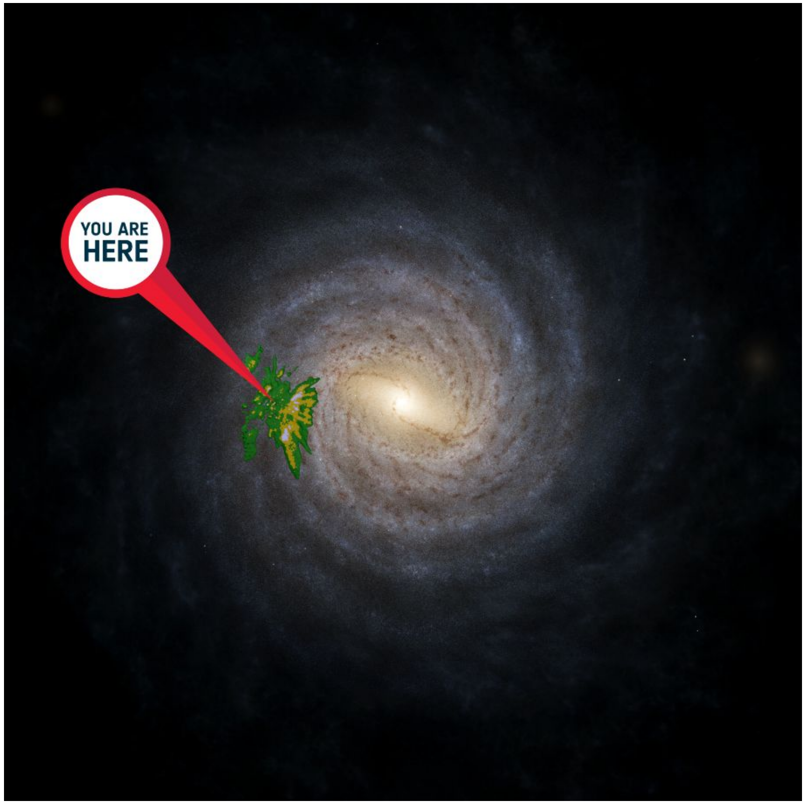

Data from Gaia

Gaia Data Release 3 (DR3), June'22:

- 1.81 billion sources
- 1.46 billion with 5-D coordinates
- 33.8 million with 6-D
- Astrophysical parameters
- Spectra:
  219 million medium-resolution
  1 million high-resolution

CURRENT DATE AND TIME	2022-12-18T18:41:54 (TCB)
MISSION STATUS	
Satellite distance from Earth (in km)	1,432,557
Number of days having passed since 25 July 2014	3068
Number of days in mission extension	1251
OPERATIONS DATA (collected since 2014/07/25)	
Volume of science data collected (in GB)	113,015
Number of object transits through the focal plane	213,827,291,050
Number of astrometric CCD measurements	2,107,726,154,629
Number of photometric CCD measurements	424,293,929,504
Number of spectroscopic CCD measurements	41,772,800,481
Number of object transits through the RVS instrument	14,029,936,285

	# sources in Gaia DR3
Total number of sources	**1,811,709,771**
	Gaia Early Data Release 3
Number of sources with full astrometry	1,467,744,818
Number of 5-parameter sources	585,416,709
Number of 6-parameter sources	882,328,109
Number of 2-parameter sources	343,964,953
Gaia-CRF sources	1,614,173
Sources with mean G magnitude	1,806,254,432
Sources with mean G_{BP}-band photometry	1,542,033,472
Sources with mean G_{RP}-band photometry	1,554,997,939
	New in Gaia Data Release 3
Sources with radial velocities	33,812,183
Sources with mean G_{RVS}-band magnitudes	32,232,187
Sources with rotational velocities	3,524,677
Mean BP/RP spectra	219,197,643
Mean RVS spectra	999,645
Variable-source analysis	10,509,536
Variability types (supervised machine learning)	24
Supervised machine-learning classification for variables	9,976,881
Specific Object Studies – Cepheids	15,021
Specific Object Studies – Compact companions	6,306
Specific Object Studies – Eclipsing binaries	2,184,477
Specific Object Studies – Long-period variables	1,720,588
Specific Object Studies – Microlensing events	363
Specific Object Studies – Planetary transits	214
Specific Object Studies – RR Lyrae stars	271,779
Specific Object Studies – Short-timescale variables	471,679
Specific Object Studies – Solar-like rotational modulation variables	474,026
Specific Object Studies – Upper-main-sequence oscillators	54,476
Specific Object Studies – Active galactic nuclei	872,228
Photometrically-variable sources with radial-velocity time series	1,898
Sources with object classifications	1,590,760,469

Astrophysics on the Cloud

Our goal: **Understand** the **history** of our Galaxy (Milky Way) and its neighbours.

Drawback: For human scales, the **timescales** of the Galaxy are **immense**, we can not do experiments. *Astronomy is an **observational** science…*

How do we do our research? We compute

- **Good** (*large*) simulations of galaxies
- with **different** (*a lot*) **parameters** to

compare with the data.

Computational time	=	Resolution of the simulation (# of particles)	×	Parameters (# of simulations)

OCRE and GalacticRainCloudS

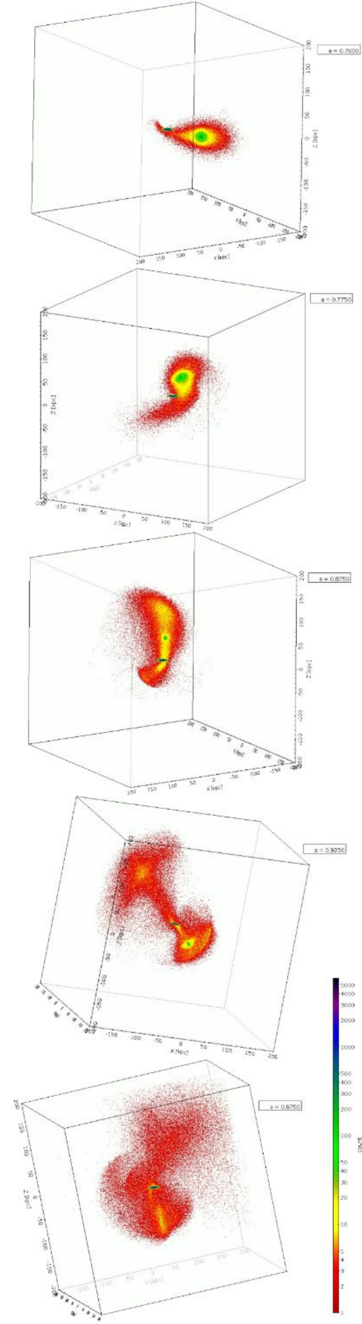

- OCRE: Open Clouds for Research Environments
  - Aimed at the European research community
- Our proposal (based on Gaia data):
  - Galactic Research in Cloud Services
  - N-body simulations to study the **Magellanic clouds** interactions
  - Study of the **kinematic substructure** of our Galaxy in 6-D, including particles simulations
  - Determine the **Initial Mass Function** of stellar formation, and the star formation history of our Galaxy
  - Other additional projects (always on Gaia data): flexibility
- Requested **100k€** on Cloud Services + support + workshops
  - Spark cluster, Linux VMs, ML service, storage…
  - Granted! →Telefónica/Altostratus + **Google Compute Cloud**
  - Started April'22, finishing next week!

Cloud functionalities used in our projects

- **Virtual Machines** (VMs), that is, **IaaS**
  - Many of them, large ones, latest architectures, long runs…
- **Storage**
  - Disks (attached to VMs) and "buckets" (shared/accessed by VMs and other services)
- Apache Hadoop/Spark **Cluster**
  - Simple configuration and use (vs. bare metal hardware)
- **Massive and high-performance database** (BigQuery)
  - Ingestion of GaiaDR3 (just GaiaSource, the main table) and some test queries
- **AI/ML services** (Vertex AI, AutoML):
  - Training (and application) of classification + regression models on Gaia data
- Data analysis tools:
  - Jupyter **notebooks**, plugged into the VMs or the cluster
- **Support**:
  - Intro to the Cloud, technical support, configurations, recommendations…
  - Tailored courses and workshops
- In general: many computational resources (quite specific processing case)
  - Not much use of powerful services like Vertex AI (lack of knowledge and expertise)

Basic use: our super-node on the Cloud

Awklusters project:

Blind all-sky search of stellar clusters,
looking for overdensities in a 5-D parameters space
(coordinates, proper motions, parallax)
incl. automatic generation of plots, statistics and summary “catalogue”.

Toy model: Just scripts on CSV files! (1.3K lines **Bash** + 1.4K lines **Awk**)

Resources: **1 large VM** (64 cores, **864GB RAM**, 2TB SSD)
Just download Gaia DR3 CSVs + Git clone + configure + launch.

Why Cloud and not a simple PC?

- Shorter run time: higher concurrency
- → Large RAM consumption (multi-dimensional Awk arrays)
- More quick tests→params. adjustment, quick re-test

Basic use: our super-node on the Cloud

Awklusters project results:

Gaia DR3 (1.8B records) processed in **~1 hour** *(~1 day in local PC)*

~1200 stellar clusters found (some may be unknown, still checking)

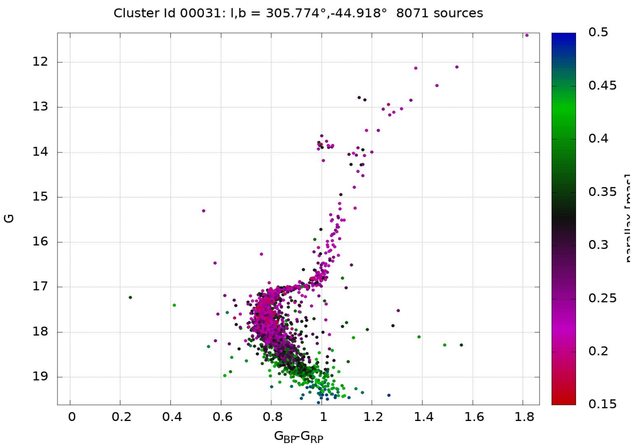

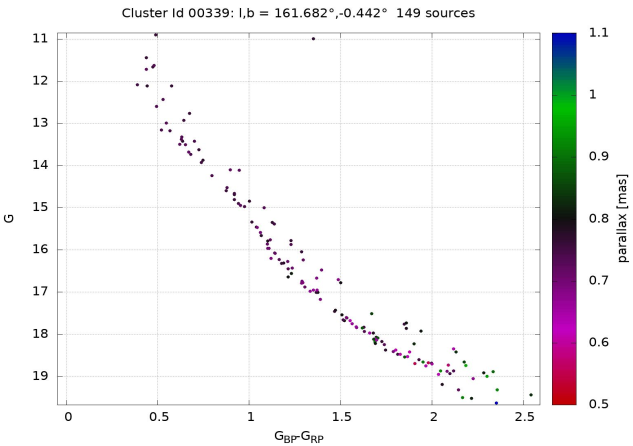

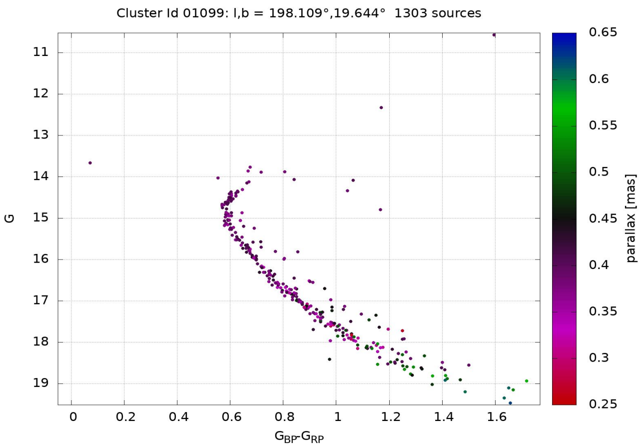

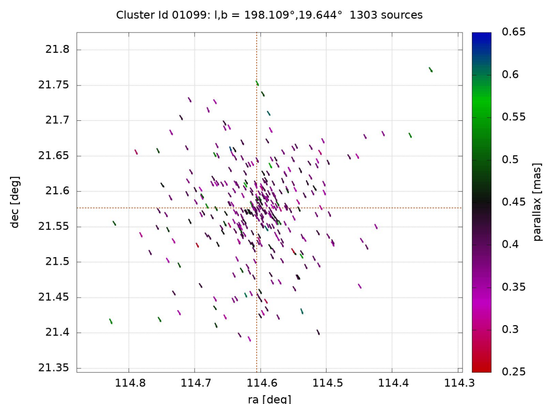

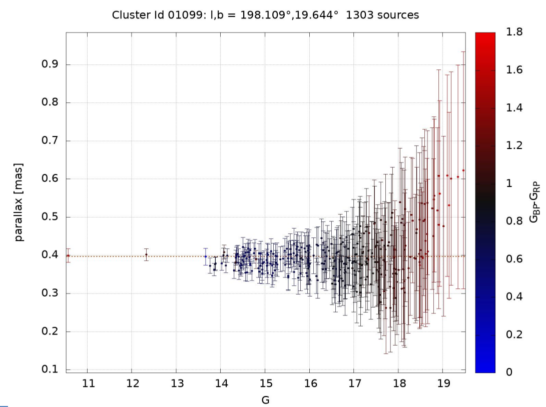

Intermediate use: our (many) nodes on the Cloud

KRATOS: a suite of **24 N-Body simulations** of the **Magellanic System**

- Large Magellanic Cloud (LMC)
- Small Magellanic Cloud (SMC)
- Milky Way (MW)

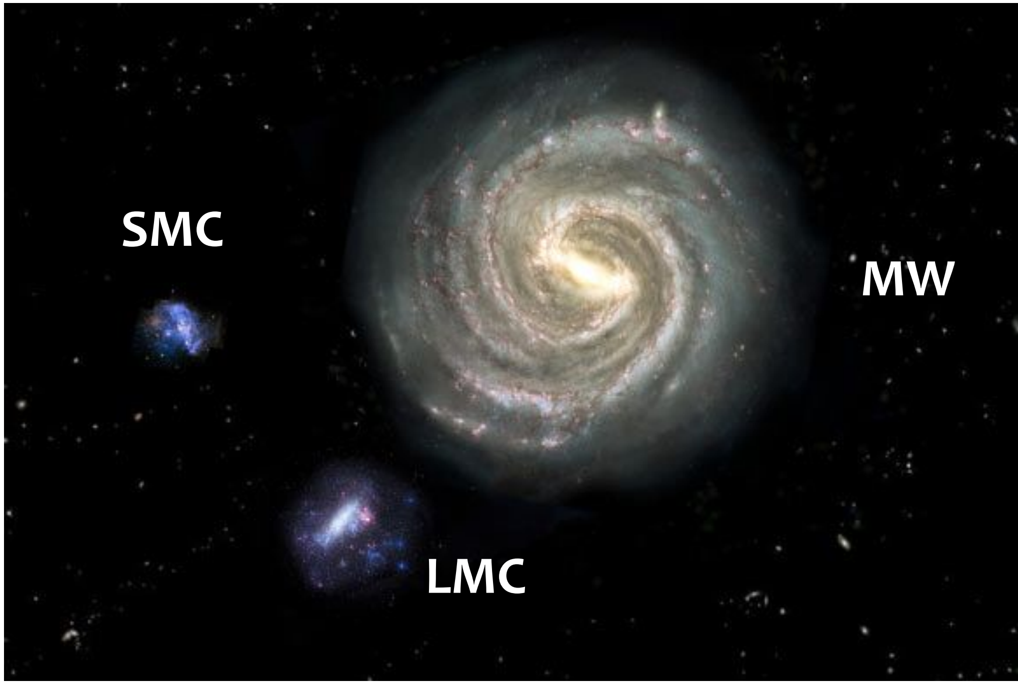

Scientific objective. Understand the properties of the LMC galaxy:

- Bar pattern speed (rotation)
- Warp
- LMC/SMC orbital story
- Magellanic bridge

Intermediate use: our (many) nodes on the Cloud

What did we need?

- Testing and setup:
  One rather large computer (16 cores, 128GB RAM, 250GB SSD),
  running **reliably** during several weeks.

- Operations:
  24 nodes (each 16 cores, 128GB RAM, 250GB SSD)
  to run **independent** simulations.

  Three weeks each simulation, **simultaneously**!

(***local machine -> ~1.4 years***)

Intermediate use: our (many) nodes on the Cloud

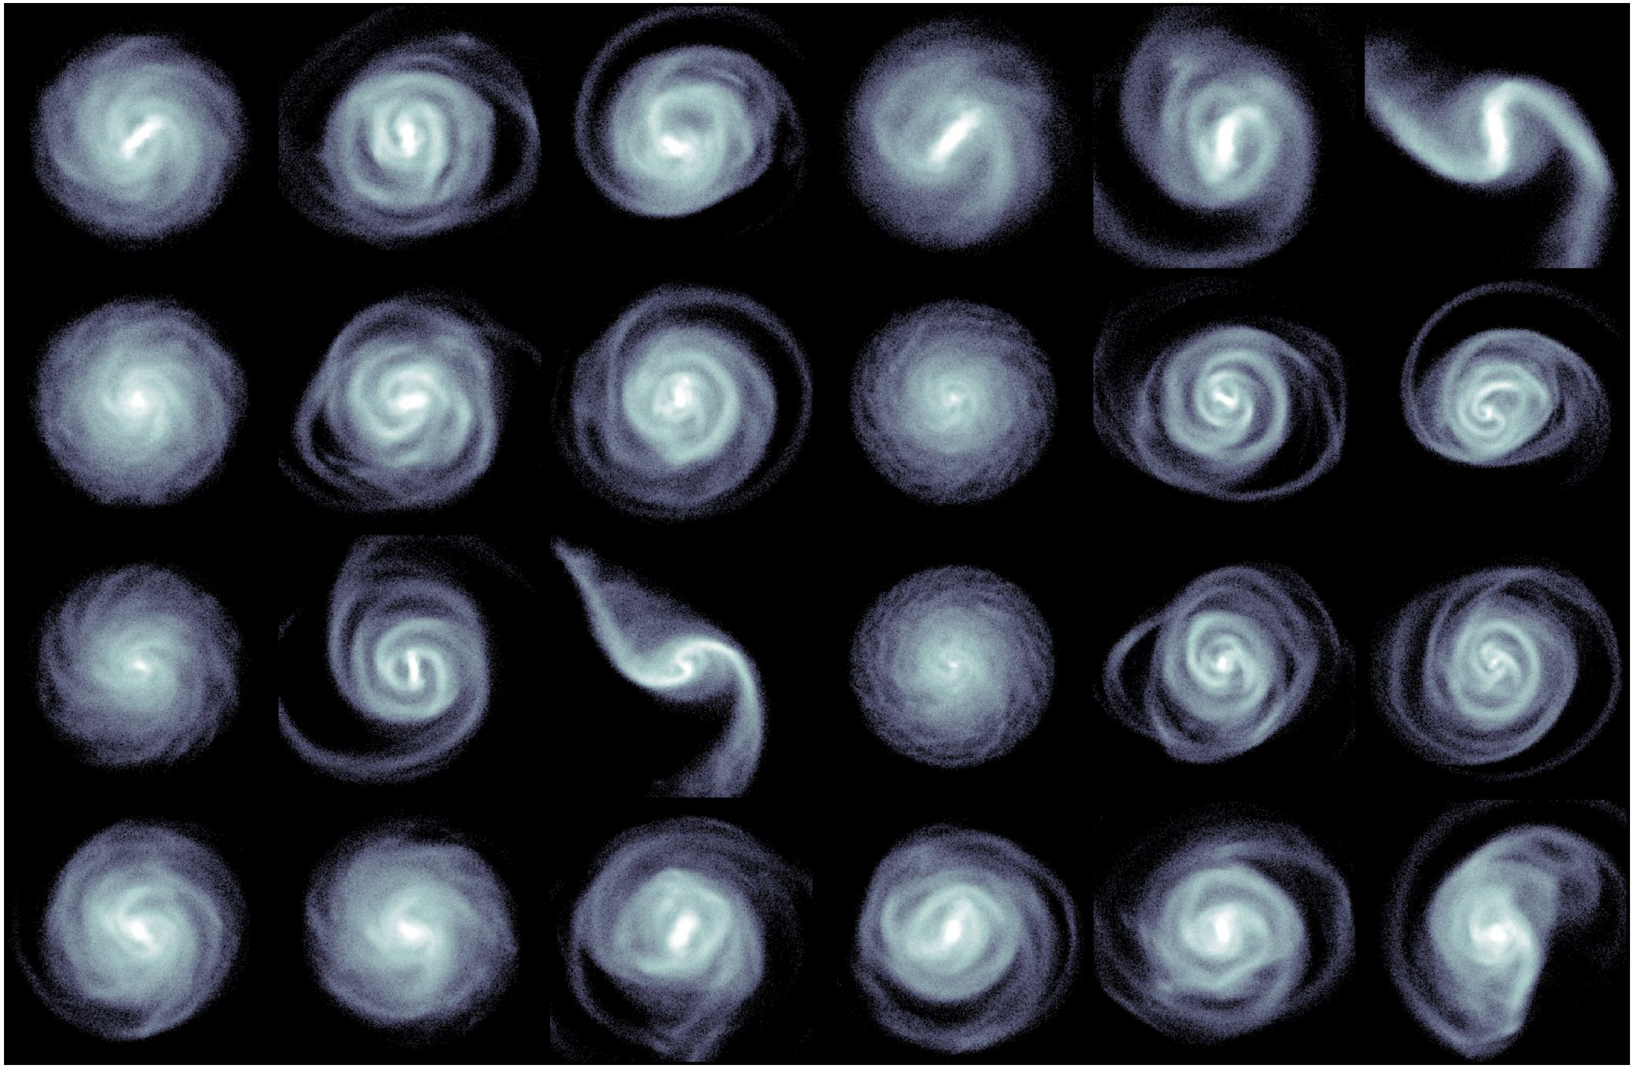

Intermediate use: notebooks on the Cloud

Disk Structure project:

Simple **analytic simulations** of a small region of the Milky Way.

Scientific objective. Understand the parameters of the bar perturbation.

4 parameters, 50 steps/parameter → 50^4 = **6.25M simulations**!

Intermediate use: notebooks on the Cloud

What did we need?

- Testing and setup:
  One rather large computer (48 cores, 48GB RAM, 500GB SSD), accessible via Jupyter Notebook.

- Operations:
  5 nodes (48 cores, 48GB RAM, 500GB SSD) to run **independent** simulations.

  One week of processing, **1.25M simulations/node**.

(***local machine -> ~4 months***)

Intermediate use: notebooks on the Cloud

Notebook setup:

- `Create VM`
- `Assign static external IP (e.g. 35.120.40.29)`
- `Add firewall rule to allow external connection to a port (e.g. 8889)`
- `Install Jupyter and dependencies.`
- `Configure Notebook (jupyter notebook --generate-config) including the previous port.`

As user:

- `Launch VM and Jupyter (e.g. jupyter notebook --no-browser --port=8889)`
- `Connect to fix IP & port (e.g. http://35.120.40.29:8889)`

Intermediate use: notebooks on the Cloud

DATA

SIMULATIONS

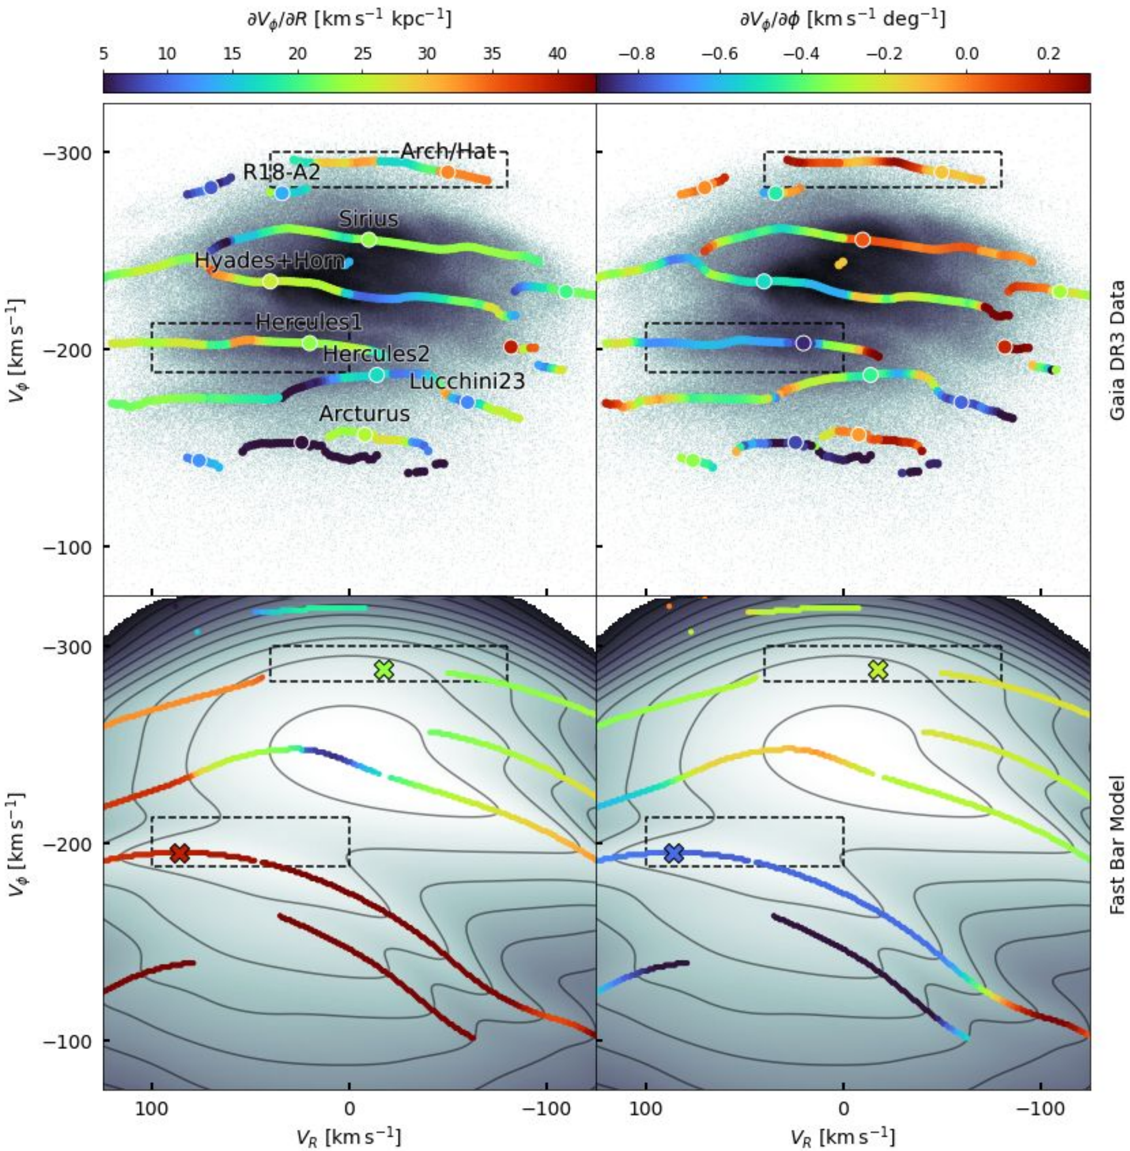

Intermediate use: small HPC cluster on the Cloud

Galactic Models project:

N-body simulations of Milky Way + Sagittarius dwarf galaxy

→study Sgr mass loss problem and phase-mixing processes in the MW

→~1 million disc particles

Technically:

- Few large nodes (128 cores, up to 850GB RAM, 10TB SSD)
- Models run with massively parallel code AREPO taking full advantage of **MPI** and uses 100% of the CPU capability

Intermediate use: small HPC cluster on the Cloud

Galactic Models project - some results:

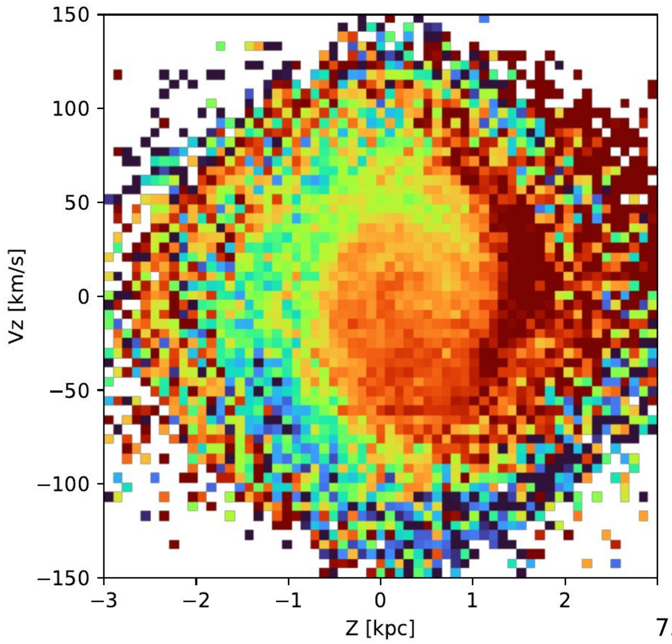

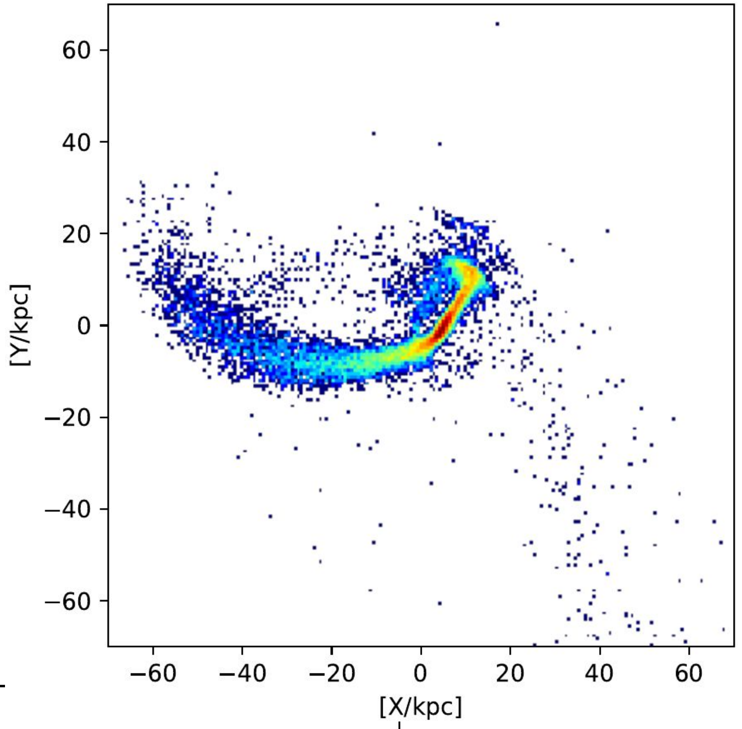

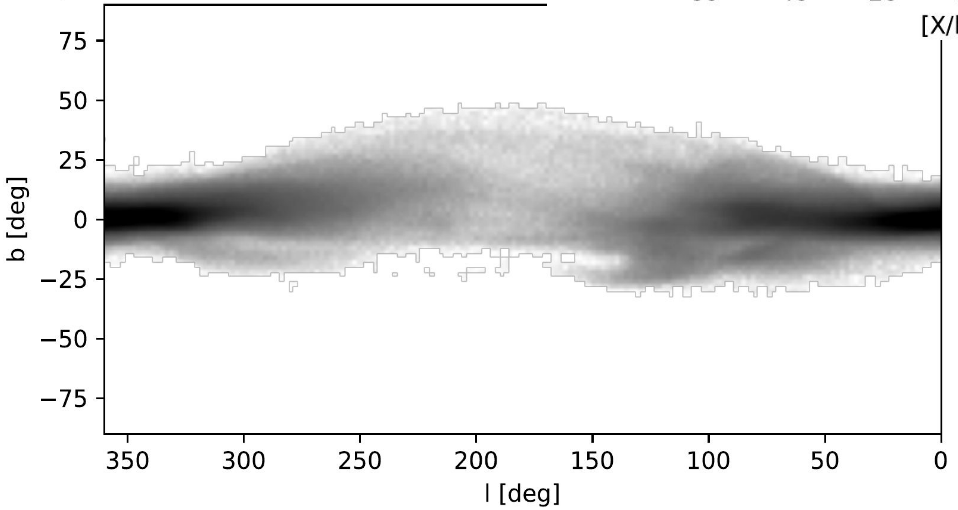

Advanced use: our Spark cluster on the Cloud

BGMFast project

Address **fundamental questions of the**
Milky Way structure and evolution in the Gaia era.

We use **data from Gaia DR3**:
magnitudes, colours, and parallaxes for stars with $G < 13$ (**7M sources**)
to explore a parameter space with **14 dimensions** (3+11)
that simultaneously includes the initial mass function (IMF) and a
non-parametric star formation history (SFH) for the Galactic disc.

Inference performed by combining the Besançon Galaxy Model fast approximate simulations (BGM FASt) and an approximate Bayesian computation algorithm.

See: BGM FASt; Mor et al. 2019 (A&A)

Advanced use: our Spark cluster on the Cloud

BGMFast project - What did we need?

Apache Spark cluster, with 1 master node and at least ~4 worker nodes (each 64 cores, 240GB RAM, 500GB disk)

Later on: **2nd independent cluster**

Very easy to setup when using cloud provider tools:
just install packages and libraries required by our software - ready to go!

Started with **notebooks** to test the infrastructure and its capabilities, migrated to **scripts** for longer executions

Code was prepared to use Spark parallelization capabilities:
Adjusting the number of parallel executions and the resources allocated to each of them was key to achieve better **performance** results

Advanced use: our Spark cluster on the Cloud

BGMFast project: Spark Application Master

▾ Active Jobs (1)

Job Id	Description ▴	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
86530	foreach at /opt/conda/default/lib/python3.7/site-packages/bgmfast/bgmfast_simulation_class.... foreach at /opt/conda/default/lib/python3.7/site-packages/bgmfast /bgmfast_simulation_class.py:477 (kill)	2023/05/27 11:35:54	0.6 s	0/1	1/40 (39 running)

▾ Completed Jobs (86530, only showing 930)

Page: 1 2 3 4 5 6 7 8 9 10 > 10 Pages. Jump to 1 . Show 100 items in a page. Go

Job Id ▾	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
86529	foreach at /opt/conda/default/lib/python3.7/site-packages/bgmfast/bgmfast_simulation_class.... foreach at /opt/conda/default/lib/python3.7/site-packages/bgmfast /bgmfast_simulation_class.py:477	2023/05/27 11:35:52	2 s	1/1	40/40
86528	foreach at /opt/conda/default/lib/python3.7/site-packages/bgmfast/bgmfast_simulation_class.... foreach at /opt/conda/default/lib/python3.7/site-packages/bgmfast /bgmfast_simulation_class.py:477	2023/05/27 11:35:50	2 s	1/1	40/40

Advanced use: our Spark cluster on the Cloud

BGMFast project: GCP cluster monitoring tools

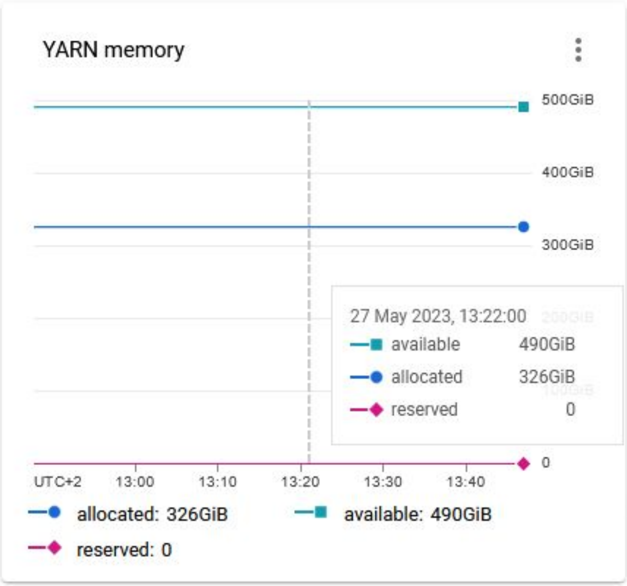

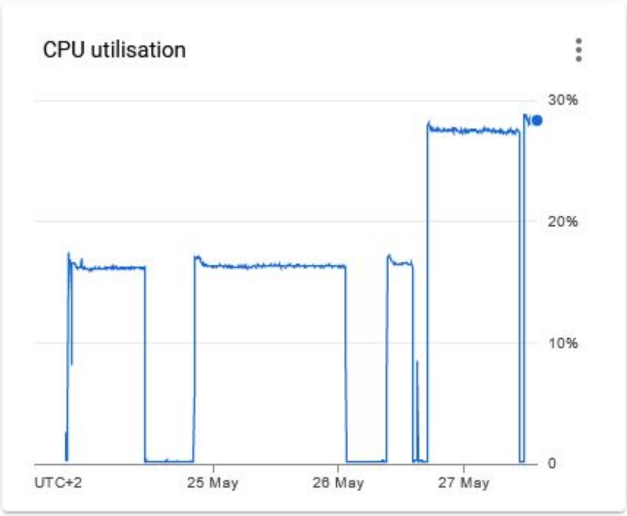

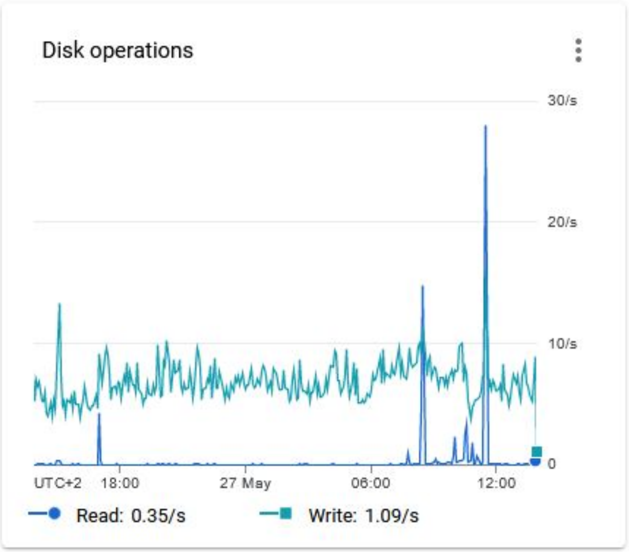

Additional Cloud services used: buckets

Buckets / Cloud storage:
Storing and retrieving a practically unlimited amount of data files.

Important to analyze how often we will access the data:

- **Standard** storage
    Frequently accessed/stored for brief periods of time.
    No fee for data retrieval; storage fee: ~0.02$ / GB / Month
        → **~40$ / Month for all GaiaSource**
- **Nearline** storage
    accessed once per month or less. Retrieval 0.01$/GB; Storage ~0.01$/GB/Month
- **Coldline** Storage
    accessed once a quarter or less. Retrieval 0.02$/GB; Storage ~0.005$/GB/Month
- **Archival** storage
    accessed less than once a year. Retrieval 0.05$/GB; Storage ~0.002$/GB/Month

Additional Cloud services used: buckets

Buckets / Cloud storage:

Different geographic areas may have slightly different prices.

Multi-region have higher availability, lower latency, but a higher cost.

Retrieval fees do not include network fees.

Beware with data ***retrieval*** fees! (“egress” data transfers)

Data can be public or private.

It allows for version control.

Additional Cloud services used: buckets

Buckets / Cloud storage:

Accessing the data: We can mount it as a local filesystem in our cloud VM

```
[user@vm ~]$ df -h
Filesystem        Size  Used Avail Use% Mounted on
/dev/sda2         500G  275G  226G  55% /
(...)
gdr3              1,0P     0  1,0P   0% /mnt/gdr3
```

Or using **gsutil**, the command line tool for interacting with cloud storage:

```
[user@vm ~]$ gsutil ls gs://gdr3/gaia_source/
```

Or using the Cloud Client Libraries:

C++, C#, Go, Java, Node.js, PHP, Python, Ruby

Additional Cloud services used: BigQuery

Tests with the **GaiaDR3 "GaiaSource" table**

1811 million records, ~1 TB

Some simple analyses in Gaia may just need (rather complex) **SQL queries**

E.g.: How many sources would I get if I apply these quality cuts?

Where are they?

Which magnitude distribution do they have? Etc.

Problem: **complex queries on 1.8 billion records!**

→ How does BigQuery perform here?

Additional Cloud services used: BigQuery

How can we (efficiently) **ingest** the data?

1. Define the dataset and the table.

   For a wild quick test: 1 **partition** every 1M sources, by "random_index"
   (learned later: probably too many partitions)

2. Define the **schema**.
   a) Automatically from CSV → difficult/impossible in our case
   b) "Manually" → actually best in our case: we already have the **SQL**

3. **Ingest.**

   In our case: **E****CSV** files → not fully supported from the Console

   From the **Cloud Shell**: **193 seconds!**

```
bq load --source_format=CSV --null_marker=null
--skip_leading_rows=1001 GDR3FullGaiaSource.GaiaSource
gs://gdr3/gaia_source/GaiaSource*.csv
```

Additional Cloud services used: BigQuery

Some quick test: **~5 seconds for this query!**

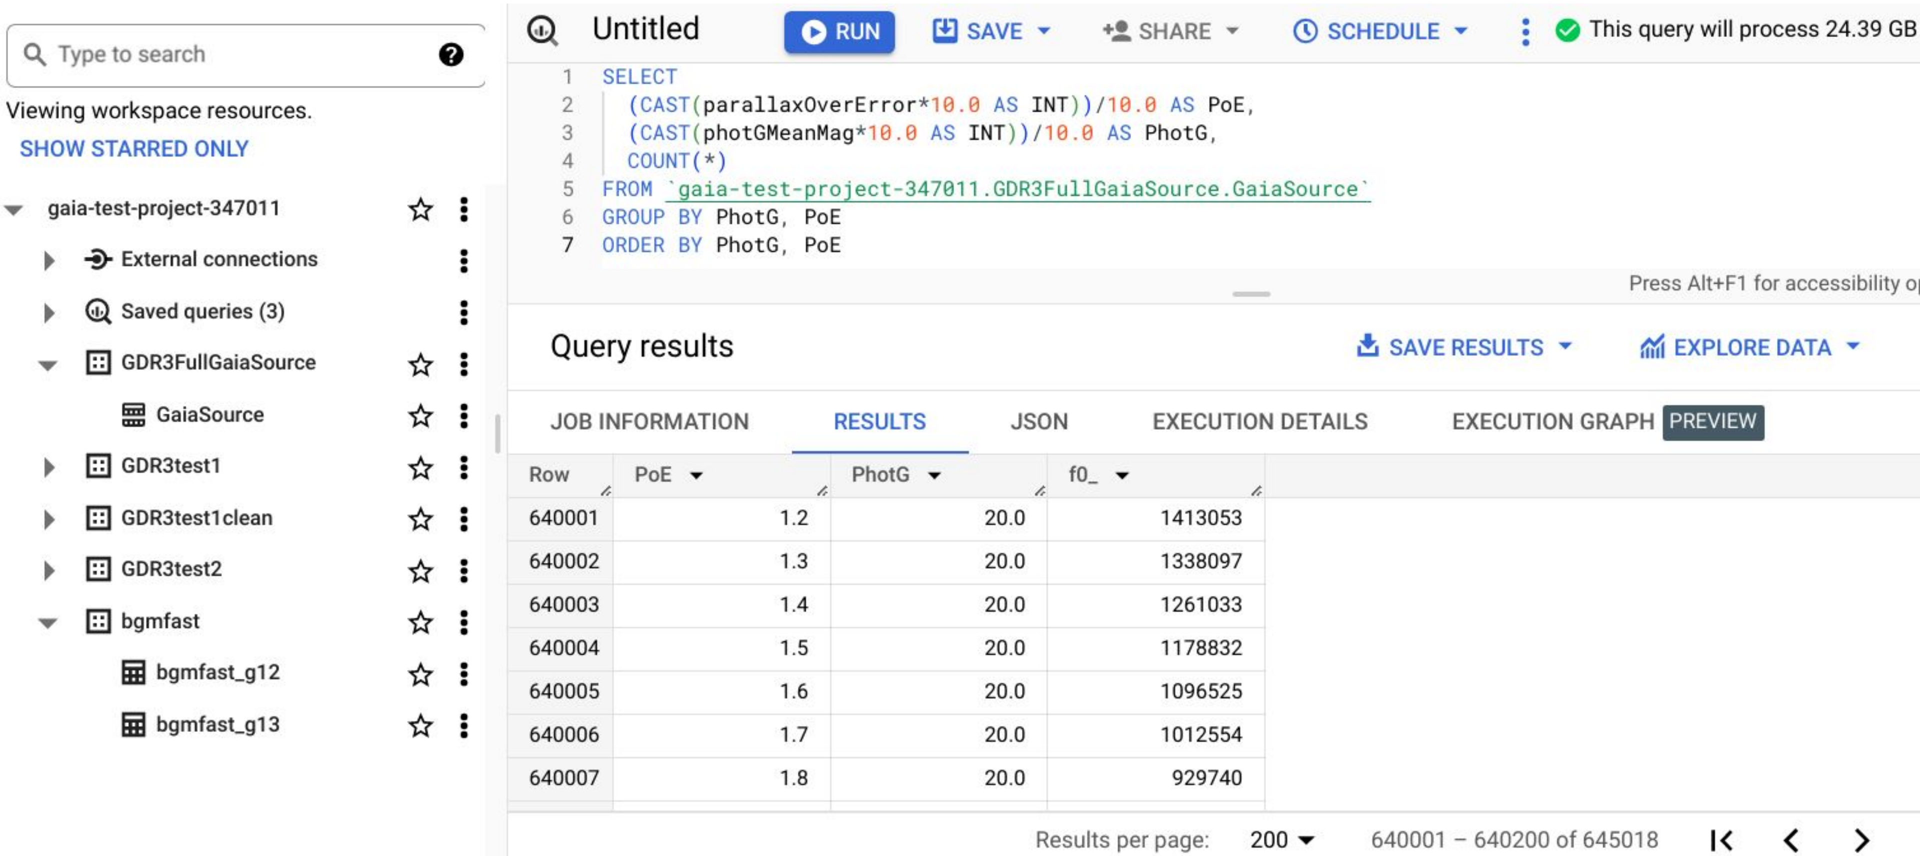

Much more complex things can be done: e.g. **ML-in-SQL**

Additional Cloud services used: BigQuery

Notebook connected to BigQuery

Plot your results!

```
[1]: %%bigquery gdr3_by_poe_and_g
SELECT
    (CAST(parallaxOverError*10.0 as INT))/10.0 as PoE,
    (CAST(photGMeanMag*10.0 as INT))/10.0 as PhotG, count(*)
FROM `gaia-test-project-347011.GDR3FullGaiaSource.GaiaSource`
GROUP BY PhotG,PoE
ORDER BY PhotG,PoE
```

Job ID ed650041-d51b-4561-ae7a-ebd52ccc6a1c successfully executed: 100%

Downloading: 100%

```
20]: gdr3_by_poe_and_g.plot.hexbin(x='PhotG',y='PoE',gridsize=181,extent=[3,21,-5,1000])
```

```
20]: <AxesSubplot:xlabel='PhotG', ylabel='PoE'>
```

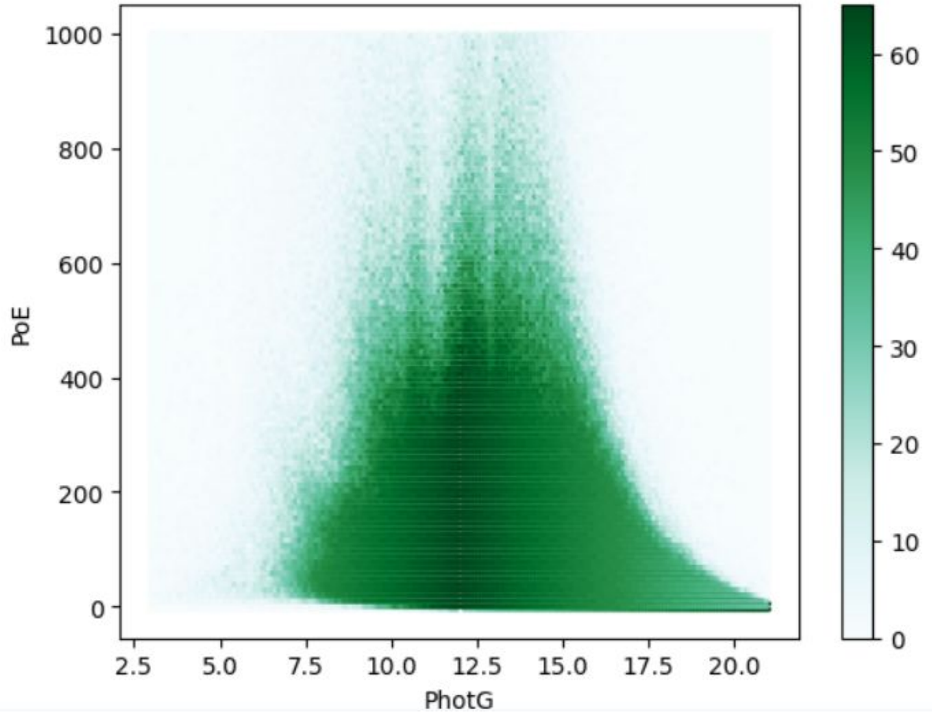

Additional Cloud services used: AutoML

VertexAI in Google → **AutoML**

Made two simple tests on a GaiaDR3 subset:

1. **Classification:** which sources are galaxies?
2. **Prediction:** what is the distance to the source?

Just did some dirty&quick test runs, few clicks here and there.

Next slides: "weights" (feature importance) automatically determined for the ~100 columns of GaiaDR3.

Additional Cloud services used: AutoML

1. **Classification:** which sources are galaxies?

Feature importance

Model feature attribution tells you how much each feature impacted model training. Attribution values are expressed as a percentage; the higher the percentage, the more strongly that feature has impacted model training. Model feature attribution is expressed using the sampled Shapley method. Learn more

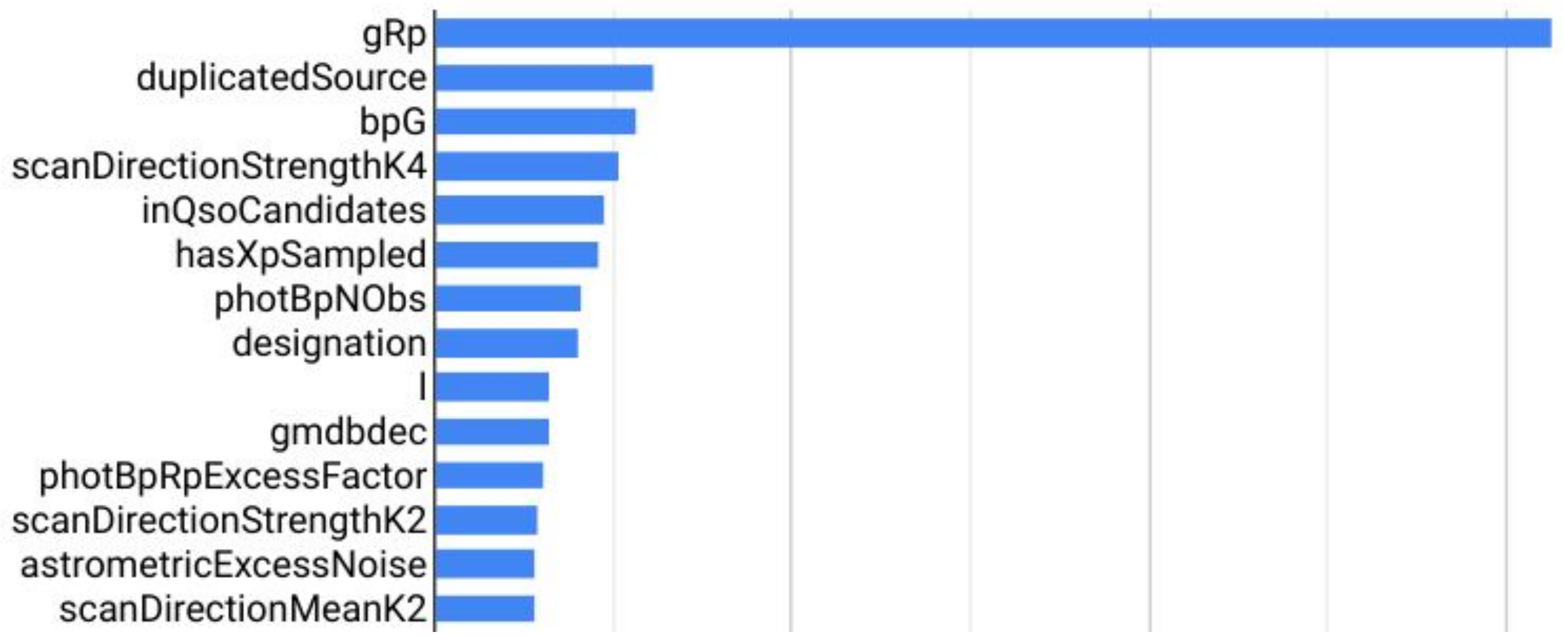

Additional Cloud services used: AutoML

2. **Prediction:** what is the distance to the source?

Target column	MAE	MAPE	RMSE	RMSLE	r^2
distanceGspphot	106.813	10.176	207.236	0.159	0.827

Feature importance

Model feature attribution tells you how much each feature impacted model training. Attribution values are expressed as a percentage; the higher the percentage, the more strongly that feature has impacted model training. Model feature attribution is expressed using the sampled Shapley method. Learn more

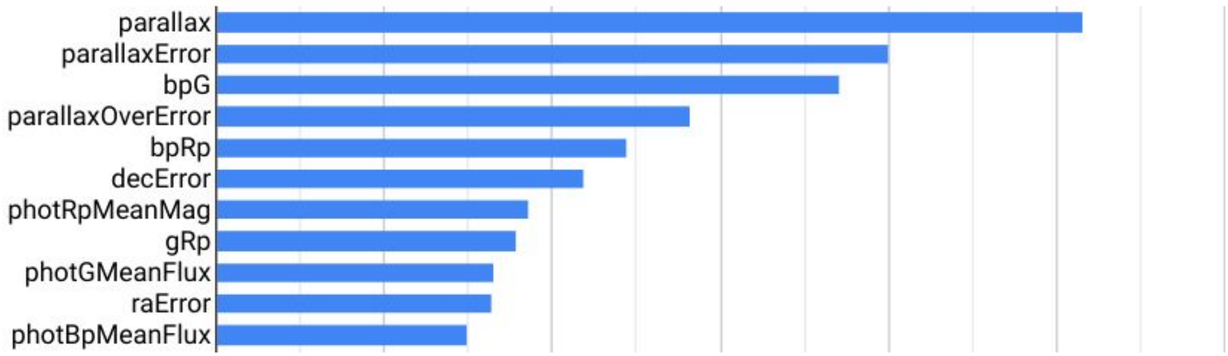

Cost monitoring

Lessons learned:

- Worth devoting a lot of effort **at the beginning** to promptly start using the resources
- Remember to adequately **label** all your resources!

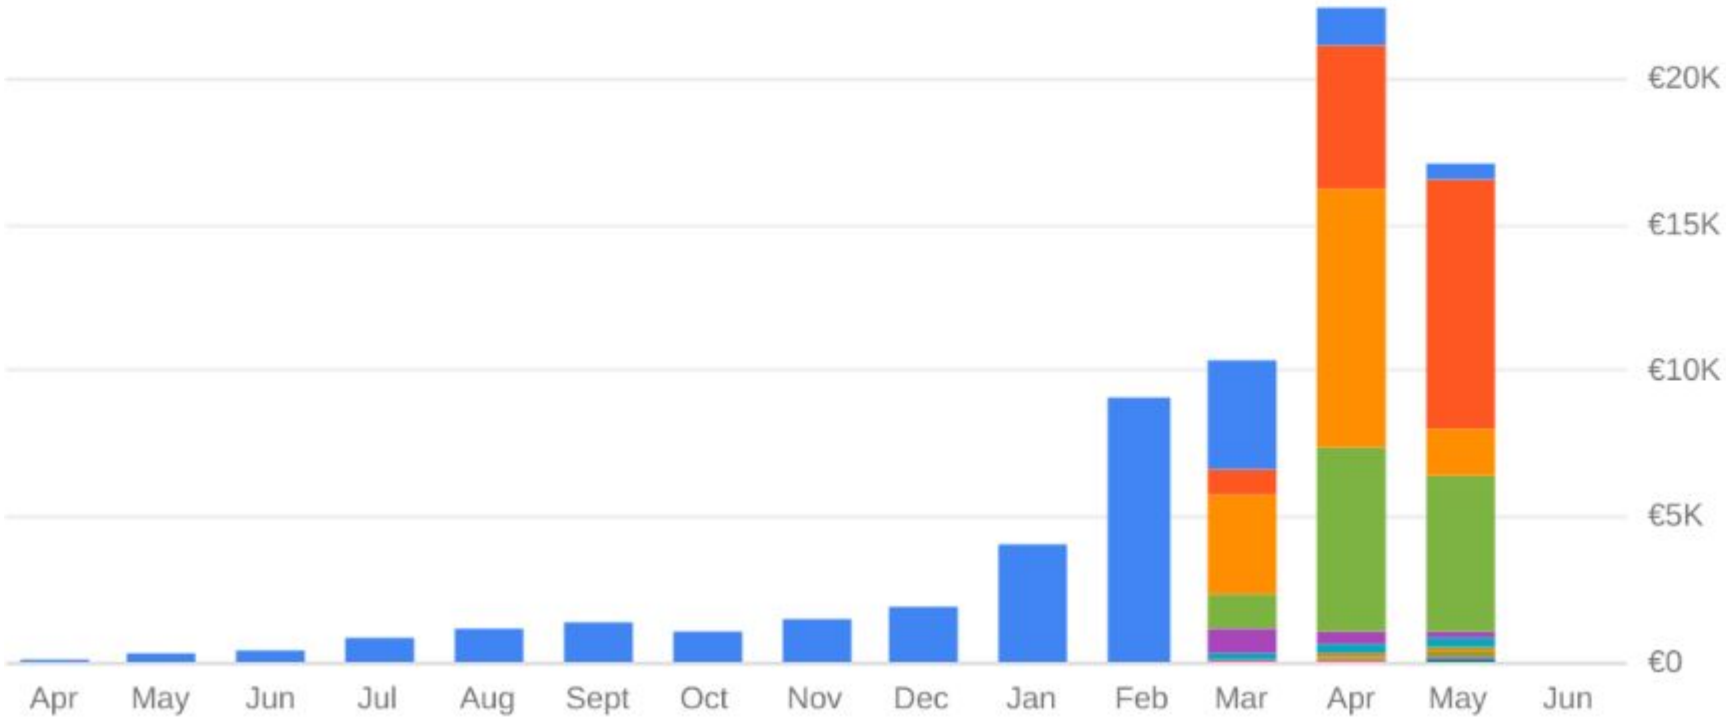

Conclusions and final considerations

- **Gaia project:** challenging number of records (not so much in disk size - for now!)
- Quite tailored data processing needs

  →**Large VMs / many VMs / Spark cluster**

  →Not far from "classical" computing approach
- More innovative/"cloudy" services used:

  Cloud storage, Big Query, ML, serverless notebooks...
- **Slow learning curve**: Most budget consumed during the last ~4 months

  →Really worth devoting a lot of time to learn and test things at the beginning!
- Good to **plan** the budget and resources usage

  Reached **>1k€/day**→be careful with costs evolution, coordinate research team
- Maybe not worth moving some processes to the Cloud…

  ...but evaluate your "**hidden costs**" of bare metal! (IT, racks, A/C, electricity...)
- Cloud is specially great for:
  - **Huge VMs** during brief periods of time
  - **Huge and efficient DBs** + analysis/visualization tools
  - **ML/AI** experiments and research
  - **Flexibility** (in our case we added some new projects very late)
  - Small groups without space/resources for bare metal

Institute of Cosmos Sciences, University of Barcelona
29-May to 2-June 2023

6th Barcelona TechnoWeek

Scientific Computing in the Cloud

Cloud use case 2: Astrophysics

Galactic RainCloudS

(Galactic Research in Cloud Services)

Jordi Portell (jportell@icc.ub.edu), Óscar Jiménez, Marcel Bernet, Sergi Bartolomé, Albert Masip, João A.S. Amarante, et al.
Universitat de Barcelona, Institut de Ciències del Cosmos (ICCUB)
Institut d'Estudis Espacials de Catalunya (IEEC)

Universitat de Barcelona
IEEC Institut d'Estudis Espacials de Catalunya
DAPCOM Data services
OCRE Open Clouds for Research Environments
esa gaia
ICCUB Institut de Ciències del Cosmos Universitat de Barcelona
Excelencia María de Maeztu 2020-2023
Unió Europea Fons Europeu Next Generation
Gobierno de España
Plan de Recuperación, Transformación y Resiliencia
Next Generation Catalunya
Generalitat de Catalunya Departament de Recerca i Universitats