Distributed SKA science-driven workflows at extreme scales: lessons from SKA precursors/pathfinders and next SKA challenges


Credits: SKA Organisation

Susana Sánchez Expósito [1], Damien Gratadour [2], Jean-Pierre Vilotte [3]

eFlows4HPC workshop: HPC Workflows for Scientific Applications

Barcelona Supercomputing Centre, 10th January 2024

(1) Instituto de Astrofísica de Andalucia – CSIC, (2) Observatoire de Paris, CNRS, (3) Institut de Physique du Globe de Paris, CNRS

Outline (1st part of the talk)

- The SKA Observatory
- The SKA Regional Centres (SRCs) and the SRC Network
- Challenges associated with the workflow execution on the SRC Network
- Workflows in the pre-SKA era: some examples.

The SKA Observatory

Open key questions in Astrophysics, Astrobiology and Fundamental Physics

- Formation of the 1st galaxies in a dark Universe dominated by atomic gas
- Evolution of the atomic gas and star formation till the current epoch
- Strong Field Tests of Gravity Using Black Holes
- Active Galactic Nuclei and the Galactic Centre
- Extrasolar planets (proto-planetary disks, biomarkers)

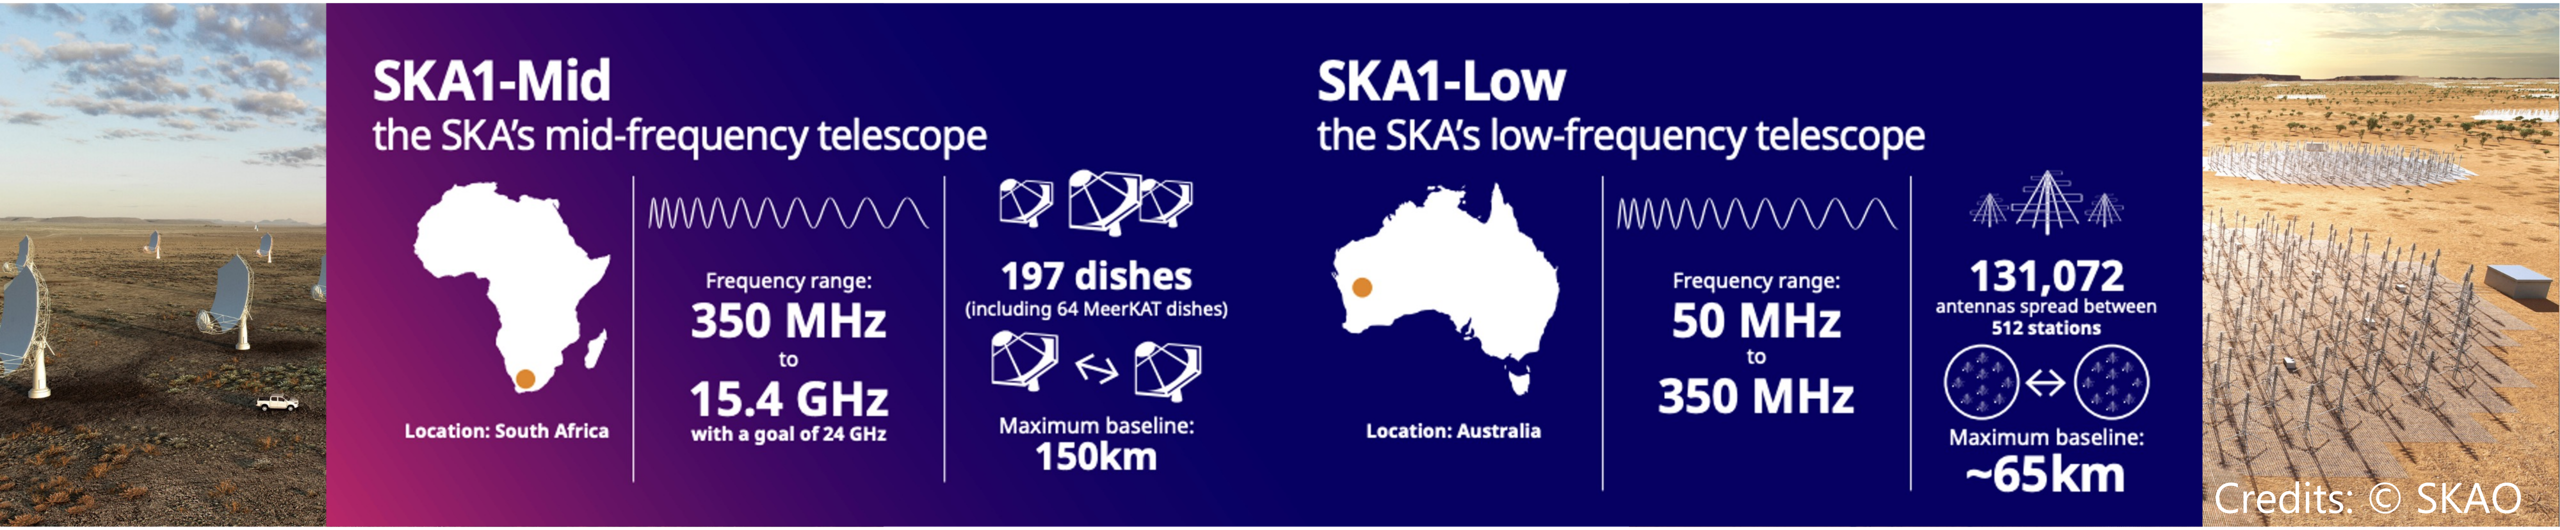

The SKA Observatory

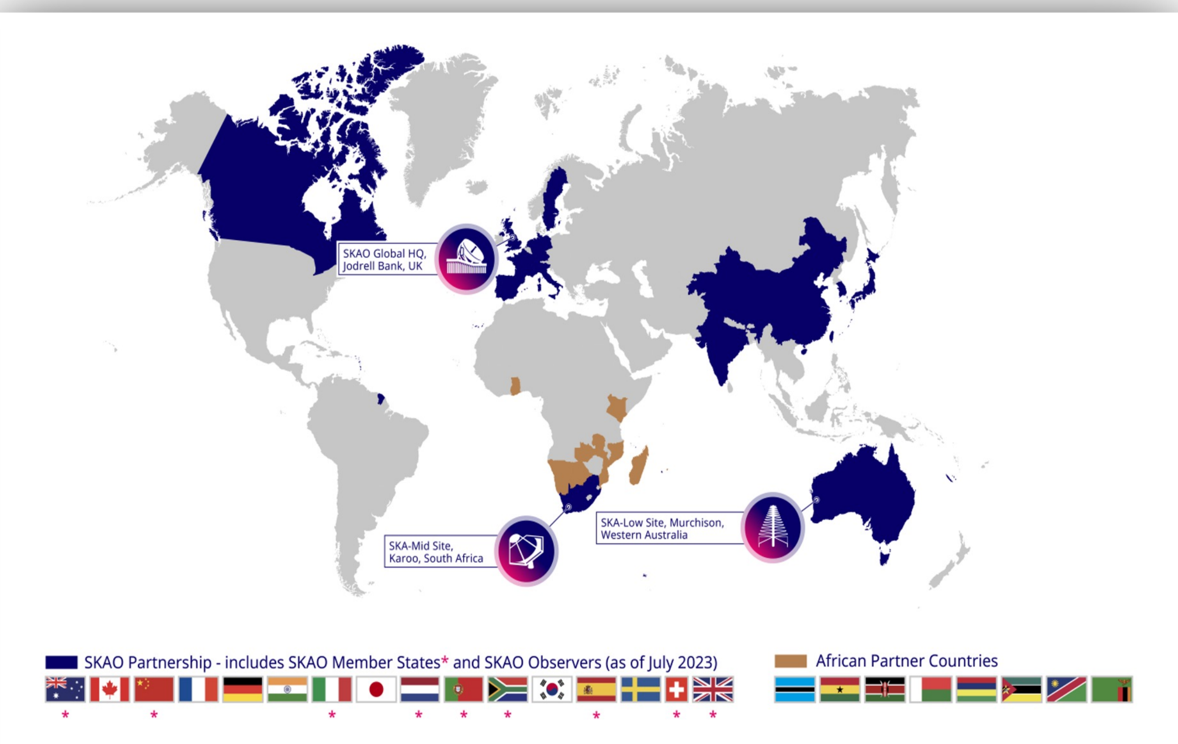

- 2013 Design consortia formed
- 2021 Construction approval
- 2024 Commissioning
- 2026 Science verification
- 2029 End of construction

Spain joined the SKA Observatory as the 9th Member country
April 2023

IAA–CSIC & SKA project:

- Coordinator of the Spanish participation in the SKA since 2011
- IAA SO programme supporting the development of the Spanish SKA Regional Centre prototype since 2019
- Participation in 3 Design consortia

  → Collab. With the BSC in the Science Data Processor consortium
- 16 IAA members participating in 7 out of the 14 SKA Science Working Groups

AMIGA group:
https://amiga.iaa.csic.es/

The SKA Regional Centre Network (SRCNet)

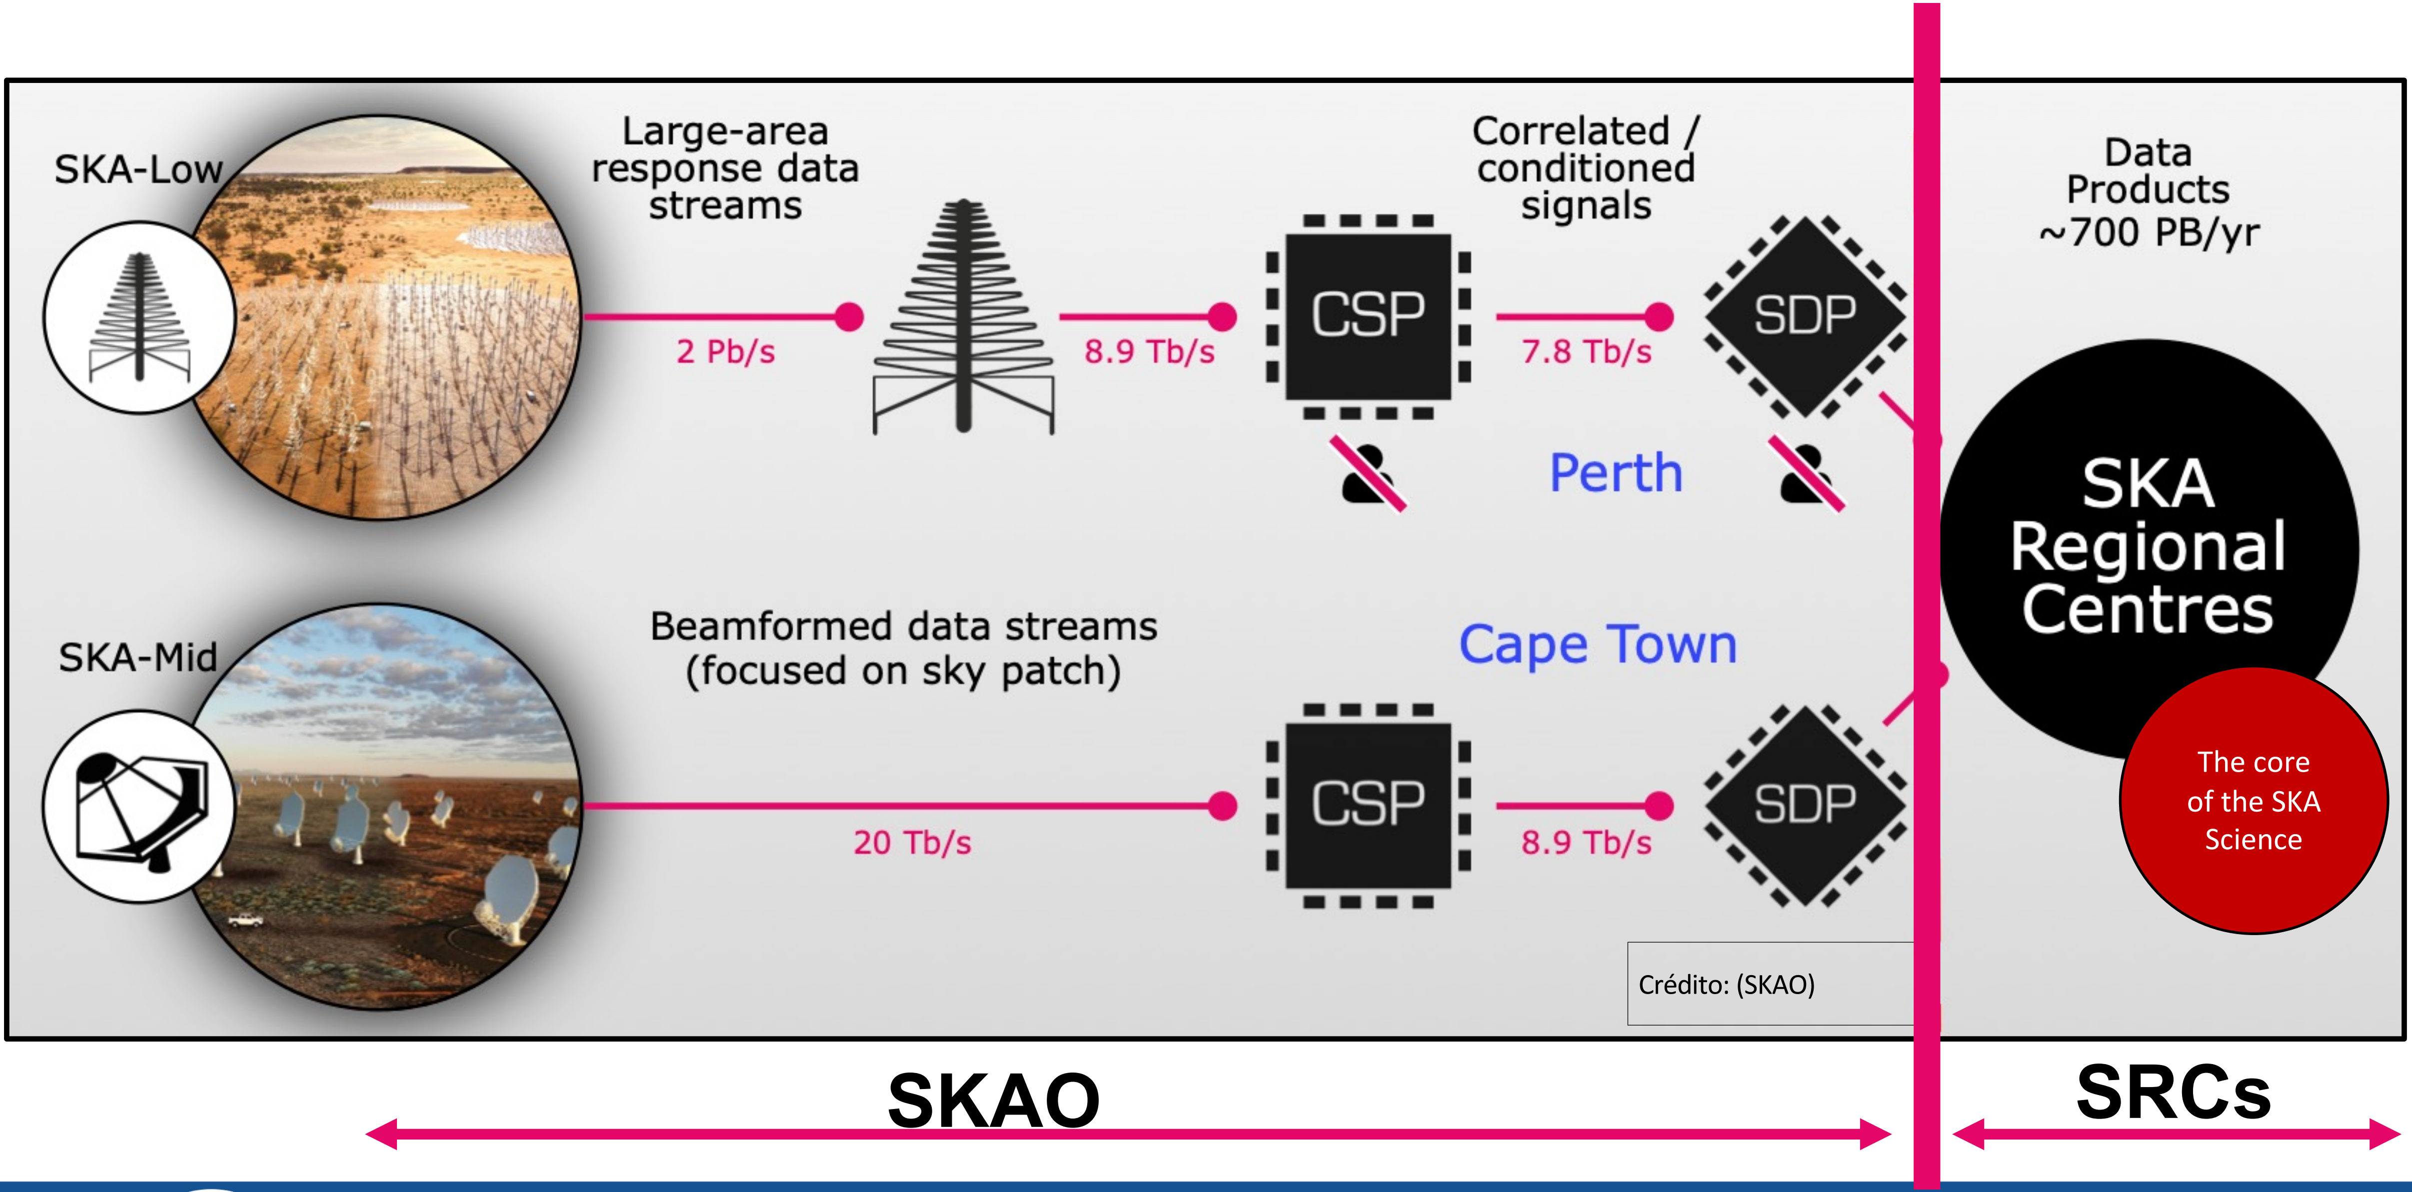

The SKA Regional Centre Network (SRCNet)

SRC Capabilities Blue print:

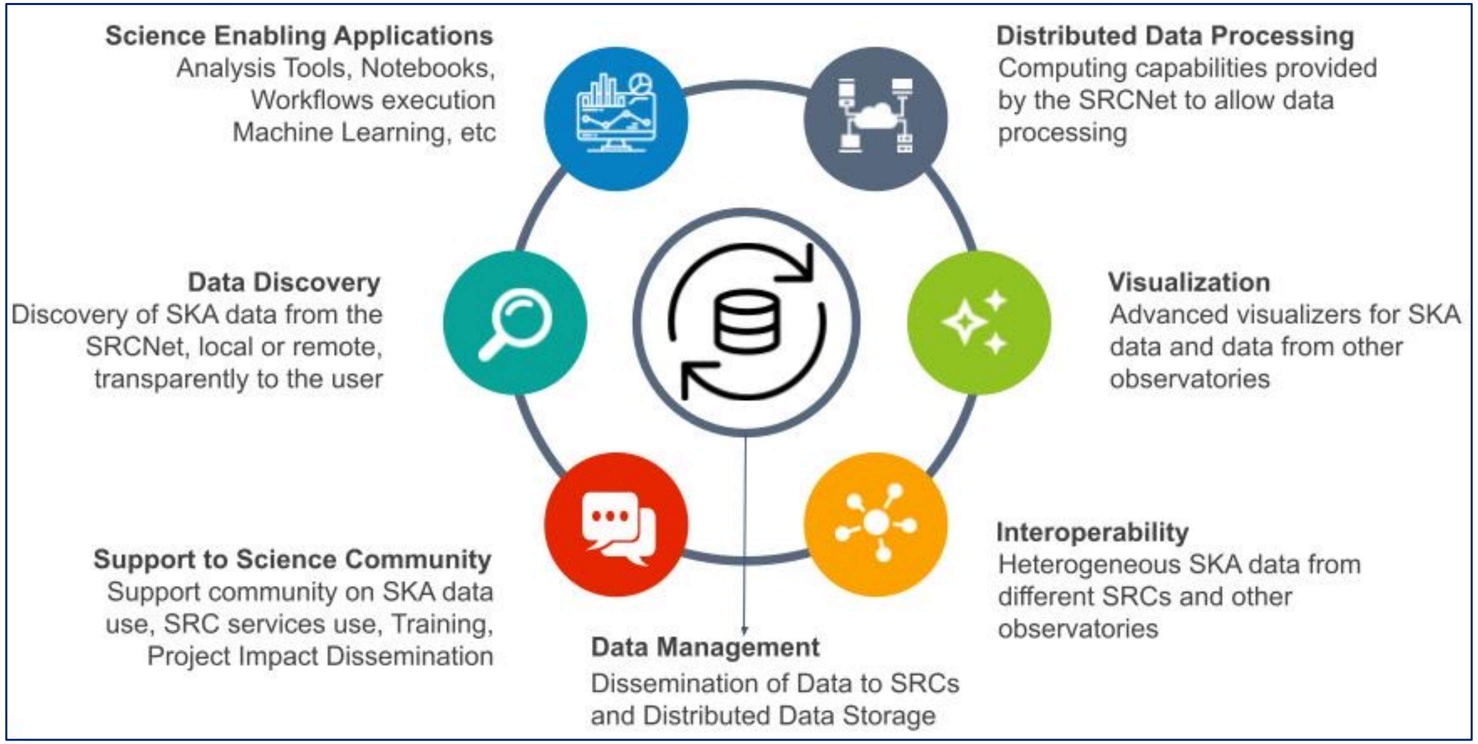

Credit: J. Salgado (SRCNet Architect)

The SRCs will

- host the SKA Science Archive
- provide access to
  - SKA data products
  - computing resources
  - scientific tools
- provide user support

SRCs as the core of the SKA Science

The SKA Regional Centre Network (SRCNet)

An ecosystem of interoperable data and services distributed in the SRCNet nodes

Credit: J. Salgado (SRCNet Architect)

- Users can access data and services, irrespective of their location or user nationality
- Same user credential to get access to data & services
- SRCNet nodes will implement common features so users are able to connect to different nodes in a transparent way
- SKA data location will be determined to optimise access and minimise data redistribution

Challenges associated with the workflow execution

Multiple nodes locally resourced and staffed, independents and heterogenous

How to bring the compute to the data and offer a seamless user experience ?

- ❑ Heterogenous **data storage solutions** → Approach based on Data Lake technology (need to align it with the diversity of datasets)

Cosmic dawn
(First stars & Galaxies)

Cosmology
(Dark matter, Large-scale structures)

Galaxy evolution
(gas content & new stars)

Cosmic magnetism
(origin & evolution)

Fundamental physics
(gravitational waves & compact objects)

Cradle of life
(Planets, Molecules, SETI)

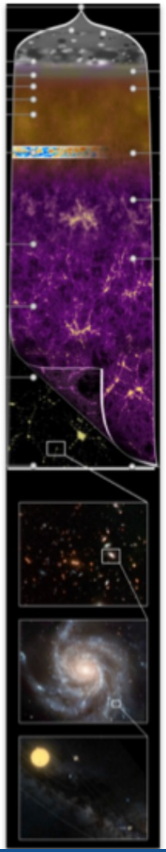

SKAO heterogenous data sets

- Image cubes
- Calibrated visibilities
- Pulsar Timing solutions
- Catalogues
- …

(3D/4D images, time series, databases, ...)

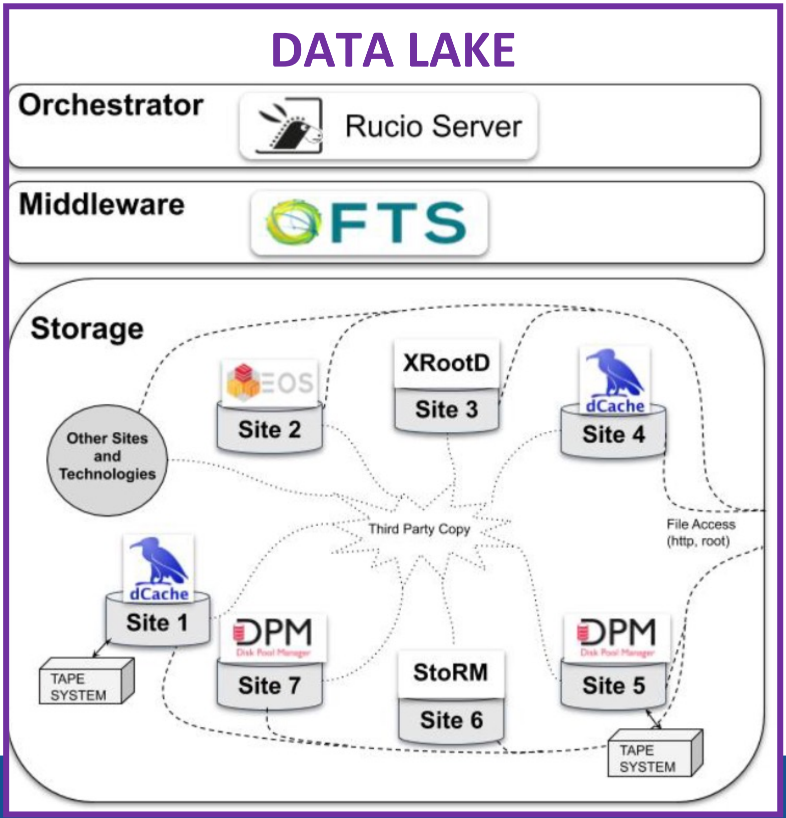

Challenges associated with the workflow execution

Multiple nodes locally resourced and staffed, independents and **heterogenous**

How to bring the compute to the data and offer a seamless user experience ?

- Heterogenous **data storage solutions** → Approach based on Data Lake technology (need to align it with the diversity of datasets)

- Heterogenous **computing framework/services** → Portability of workflows → Approach based on container technology (enough?, "Abstraction layer with portable programming models")

Challenges associated with the workflow execution

Multiple nodes locally resourced and staffed, independents and **heterogenous**

How to bring the compute to the data and offer a seamless user experience ?

- ❑ Heterogenous **data storage solutions** → Approach based on Data Lake technology (need to align it with the diversity of datasets)

- ❑ Heterogenous **computing framework/services** → Portability of workflows → Approach based on container technology (enough?, “Abstraction layer with portable programming models”)

- ❑ Nodes have "**non-uniform policies** in security, access and allocation of resources under the administration of local SRC nodes."

Challenges associated with the workflow execution

Multiple nodes locally resourced and staffed, independents and **heterogenous**

How to bring the compute to the data and offer a seamless user experience ?

- ❑ Heterogenous **data storage solutions** → Approach based on Data Lake technology (need to align it with the diversity of datasets)

- ❑ Heterogenous **computing framework/services** → Portability of workflows → Approach based on container technology (enough?, “Abstraction layer with portable programming models”)

- ❑ Nodes have "**non-uniform policies** in security, access and allocation of resources under the administration of local SRC nodes."

- ❑ Nodes will have a different set of analytical tools / **computing capacities**

  → "Centralised intelligent resource management: appropriate resource allocation knowing the behaviour/request of workflows, policies, topology and current status of the systems.”

Challenges associated with the workflow execution

❑ Centralised intelligent resource management → IVOA Execution Planner

Define protocols to identify where to execute a particular workflow

”IVOA Execution Planner”
https://wiki.ivoa.net/internal/IVOA/ExecutionPlanner20211104/IVOA-EP-note.pdf

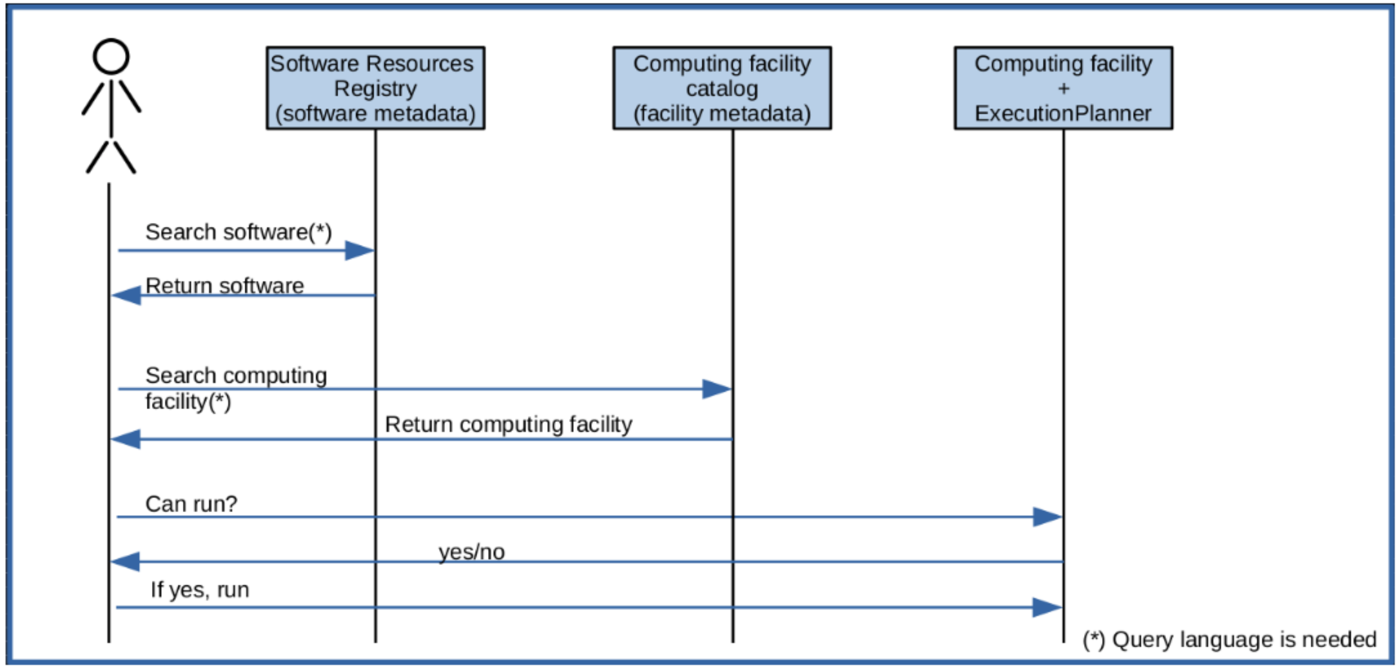

Figure 3. Use case: a user searches for software in a repository, then, having a software description, that in the simplest case is a unique identifier, searches in a registry a suitable facility where run the software. At the end he is able to run the software.

Modeling software solutions and computation facilities for FAIR access. S. Bertocco. https://arxiv.org/pdf/2302.11447.pdf

Challenges associated with the workflow execution

Complexity and **Heterogeneity** of the SKA workflows (HPC, HDA, AI, Visualisation)

- Diverse patterns of when, where and how data are accessed, transformed, analysed and intermediate results managed
  - → "Control and flexibility of data and compute placement in run time"
- "Workflows portability and composability"
- Provenance System: "Make available provenance streams captured by execution of workflows"
  - → To support FAIR principles
  - → And the "Centralised intelligent resource management"

SRCNet will embrace FAIR principles[1]

→ *Data processed within the SRC Network will automatically propagate all metadata and provenance information in support of FAIR principles*.

[1] SRCNet Vision and principles by the SRC Steering Committee

Workflows in the pre-SKA era: some examples by AMIGA group

Pathfinder telescopes

Precursor telescope

Low-Frequency Array (LOFAR)

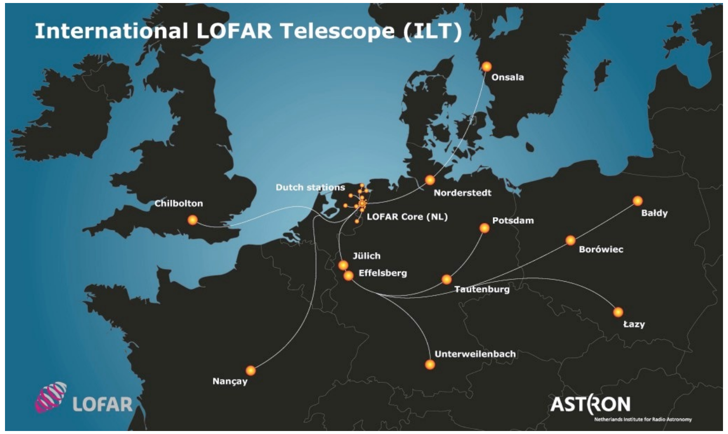

Expanded Very Large Array (EVLA)

MeerKAT

Workflows in the pre-SKA era: some examples by AMIGA group

Web services as Building Blocks for Science Gateways in Astrophysics, Sanchez-Exposito+
http://link.springer.com/article/10.1007/s10723-016-9382-y **2016**
Calibration of LOFAR data on the cloud, Sabater+.
https://doi.org/10.1016/j.ascom.2017.04.001 **2017**

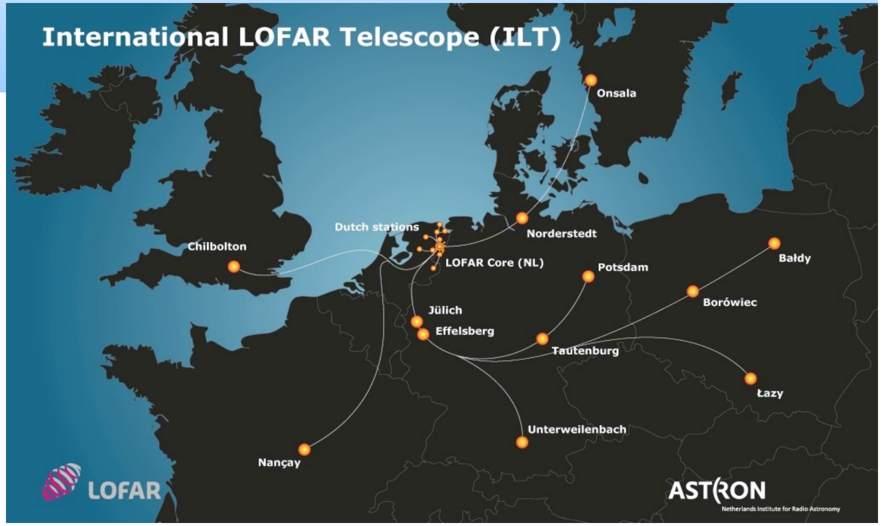

- Size of a single observation: 3-4 TBs
- Processing of chunks of data in parallel
  - → Executing 3-steps workflow for each chunk
  - → 1st step only uses 1 core

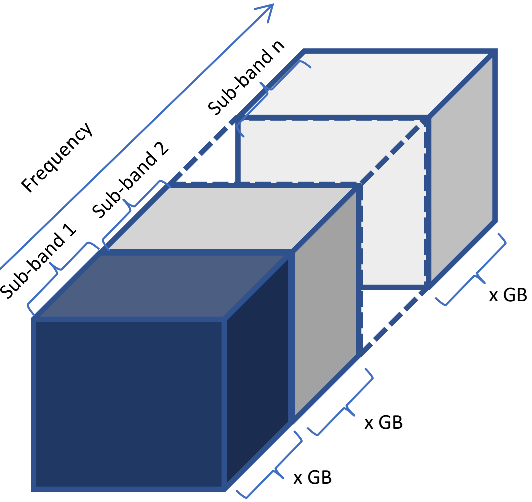

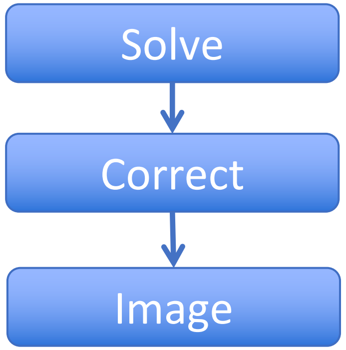

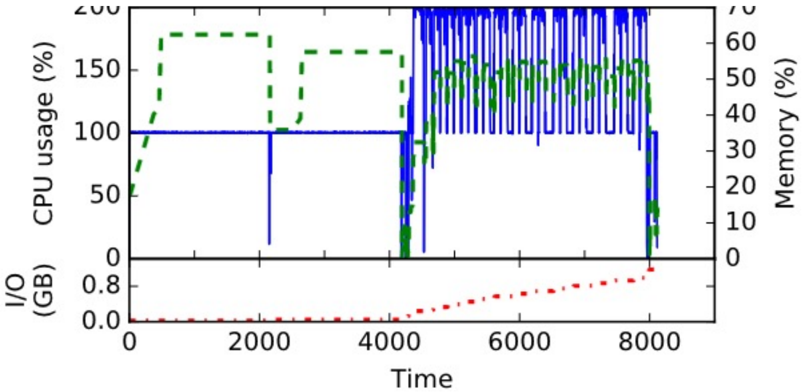

Workflows in the pre-SKA era: some examples by AMIGA group

A case study of the Hi content of Hickson Compact Group16. M.G. Jones+
https://doi.org/10.1051/0004-6361/201936349, **2019**

Enabling end-to-end reproducibility of the scientific studies
→ from the initial data processing to the plots/figures of the paper

CASA
Common Astronomy Software Applications

Flagging and Calibration

CASA
Common Astronomy Software Applications

Imaging

SoFiA
Source Finding Application

Slicer Astro

Manual step

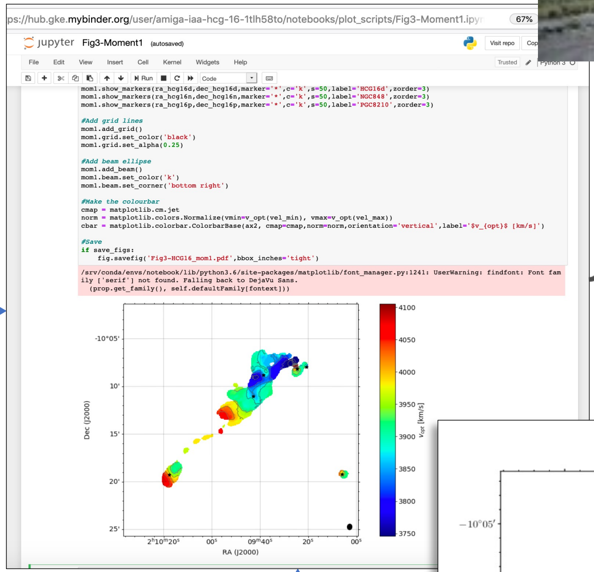

jupyter

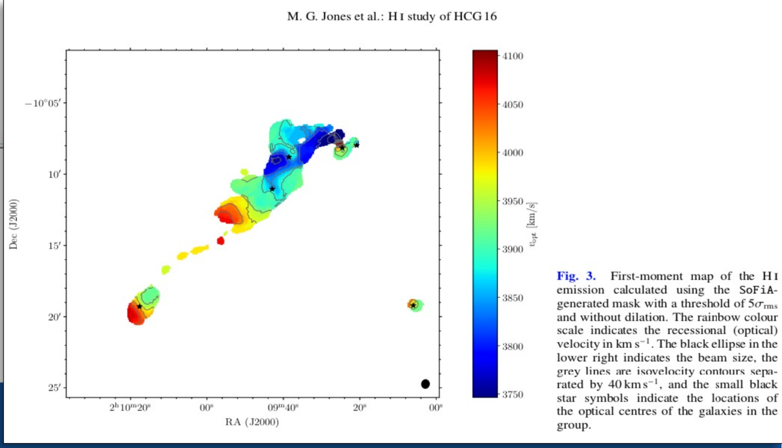

Fig. 3. First-moment map of the H I emission calculated using the SoFiA-generated mask with a threshold of $5\sigma_{rms}$ and without dilation. The rainbow colour scale indicates the recessional (optical) velocity in km s^{-1}. The black ellipse in the lower right indicates the beam size, the grey lines are isovelocity contours separated by 40 km s^{-1}, and the small black star symbols indicate the locations of the optical centres of the galaxies in the group.

Workflows in the pre-SKA era: some examples by AMIGA group

“Accelerated evolution in the densest groups of galaxies: MeerKAT imaging of the missing HI” (Observing proposal, Verdes-Montenegro+)

- 6 targets (Hickson Compact Groups)
- Raw data size: ~ 50TB (~9TB each target)
  → Selecting sub-datasets in origin to minimise data transfer
- Processing target by target (CARACAL)
  - 100 GB input data set
  - Intermediate data: 5x input data (~500GB)
- Source finding (SoFIA)
  - Final results: ~10GB

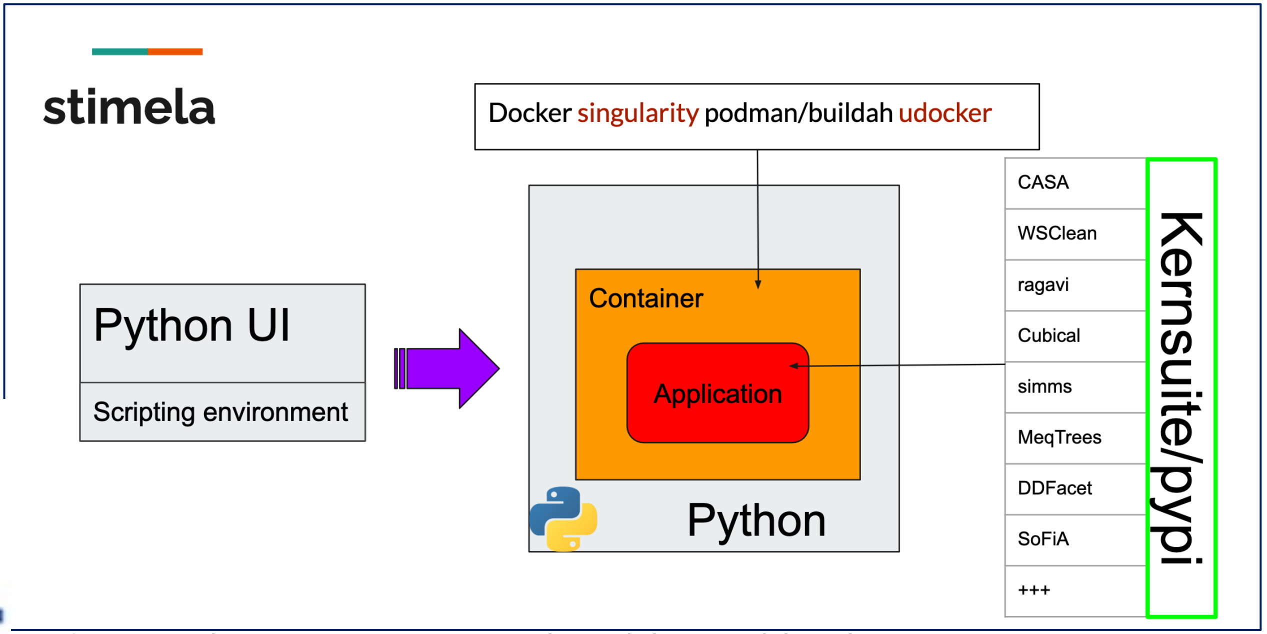

Credit: Gyula I. G. Jozsa & Sphesihle Makhathini
https://indico.in2p3.fr/event/21698/contributions/84474/

Workflows in the pre-SKA era: some examples by AMIGA group

SKA Data Challenge 2 (2021)

Multifrequency source finding and characterization of HI emitting galaxies.

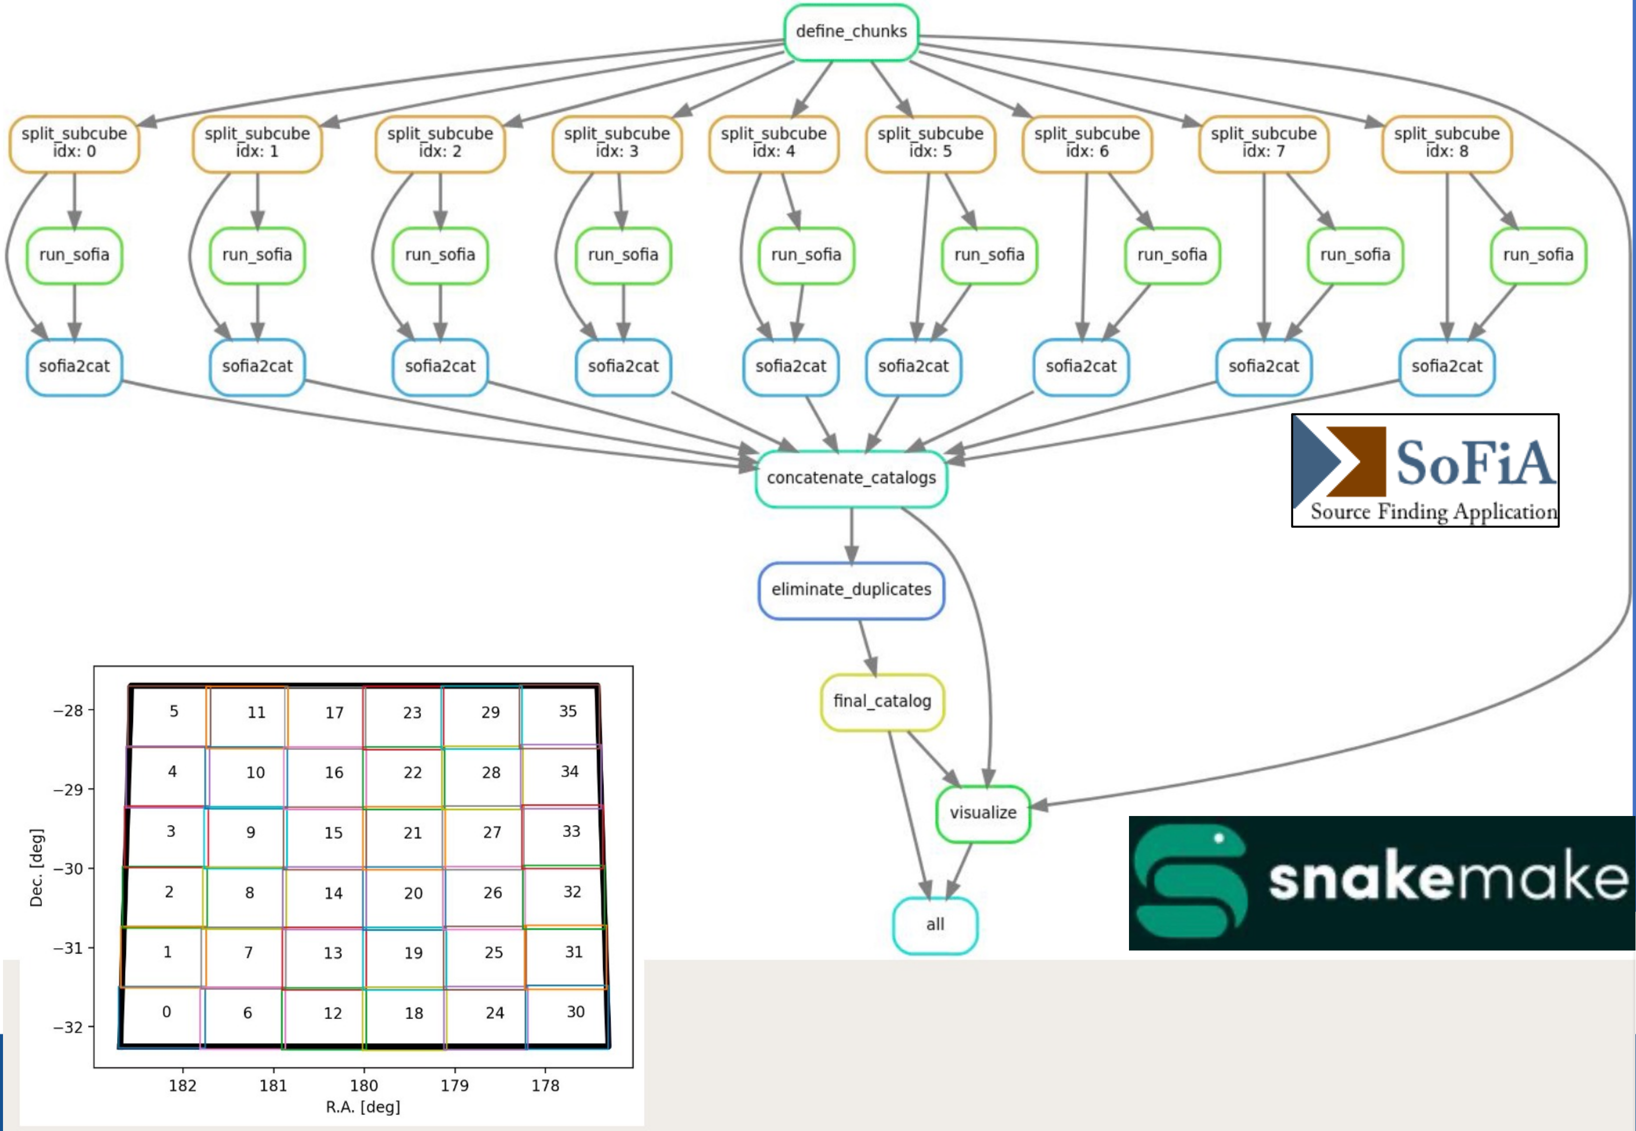

Solution by HI-Friends team

https://zenodo.org/records/6802188 (Moldon+)

- 1 TB input data set
- Data divided in overlapping chunks
- Chunks processed in parallel
- SoFiA for source finding and characterization
- Final step to eliminate duplicated sources

→ Several runs for tuning SoFiA parameters and get the best results

Workflows in the pre-SKA era: some examples by AMIGA group

SKA Science Data Challenge 2: analysis and results
https://doi.org/10.1093/mnras/stad1375, Hartley+

Position	User	Group	Score	
1	minerva	MINERVA	23254.16	**AI/ML (CNN)**
2	forska	FORSKA-Sweden	22489.43	**AI/ML (CNN)**
3	sofia	SoFiA	16822.24	SoFiA Source Finding Application
4	naoc-tianlai	NAOC-Tianlai	14416.02	SoFiA Source Finding Application
5	hi-friends	HI-FRIENDS	13902.62	SoFiA Source Finding Application
6	epfl	EPFL	8515.16	"Joint likelihood" Algorithm
7	spardha	Spardha	5614.59	SoFiA Source Finding Application
8	starmech	Starmech	2095.65	SoFiA Source Finding Application
9	jlrat	JLRAT	1079.73	**AI/ML (CNN)**
10	coin	Coin	-1.76	**AI/ML (CNN)**
11	hiraxers	HIRAXers	-2.00	"Peak finding" Algorithm
12	shao	SHAO	-471.00	SExtractor Programme

Reproducibility Award

- EPFL Bronze
- FORSKA-Sweden Silver
- HI-FRIENDS Gold
- NAOC-Tianlai Bronze
- SHAO Bronze
- Team SoFiA Silver

Conclusions

- SKAO: A Big Data instrument that will address key questions in Astrophysics, Particle Physics and Astrobiology

  → Diversity in scientific use cases

- SKAO data will be scientifically analysed in the SRC Network

  →SRC Network, the access point to the SKAO data

- The SRCNet, a network of independent and heterogenous nodes

  →Different storage solutions, computing framework, capacities and policies

- Workflows in the pre-SKA era. Use Case: HI in Galaxies

  →Several scientific applications, with different requirements, in the same workflow

  →Use of containerisation to improve portability and reusability

  →Parallelization provided by the data division in chunks

Thanks

With the financial support from a) and b):

a) MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES
Financiado por la Unión Europea NextGenerationEU
Plan de Recuperación, Transformación y Resiliencia
AGENCIA ESTATAL DE INVESTIGACIÓN

b) The grant **CEX2021-001131-S** funded by MCIN/AEI/ 10.13039/501100011033