GridLab: past, present & future

Rob van Nieuwpoort | December 2023

GridLab

Universiteit Leiden
The Netherlands

A problem we tackled in GridLab

Deployment: starting and controlling of applications and tasks, remote data access

There are many different platforms and infrastructure providers

Laptop, PC in the lab, local cluster, supercomputer, cloud, …

University, regional data center, national lab / supercomputing center, AWS, Azure, …

All with their own protocols, libraries and access methods

Slurm, GridEngine, Torque, Amazon-Batch, …

SFTP, GridFTP, WebDAV, S3, Hadoop, DataVerse, iRODS, …

SSH key, OpenID, SURFConext, Google, MS account, …

Researchers don't want to deal with this complexity

A GridLab success story

- Grid Application Toolkit (GAT), 2004
- Evolved into OGF standard SAGA, 2009
  (Simple API for Grid Applications)
- Evolved into XENON
  Built by the Netherlands eScience Center, 2014 - today

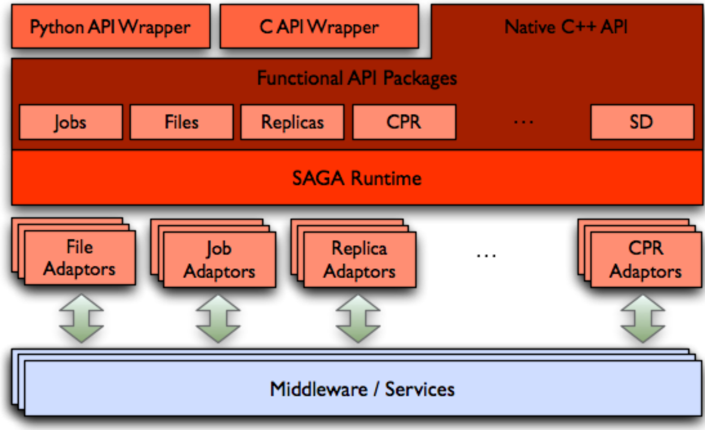

SAGA
A Simple API for Grid Applications

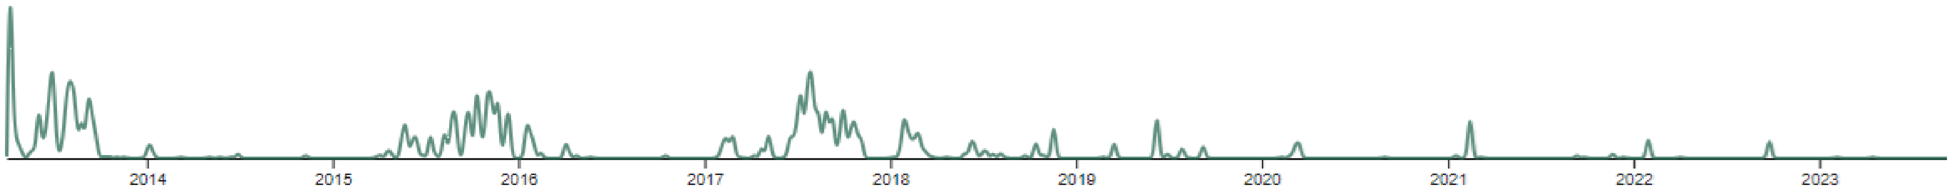

CONTACT PERSON

Jason Maassen
Netherlands eScience Center

netherlands eScience center

Xenon ecosystem

Used by 14 tools,

Which are used in many dozens of projects

Abstraction layer

- workflow frameworks CWL, snakemake, ...
- web service
- AMUSE
- KernelTuner GPU optimization
- XENON
- Noodles Python programming model
- commandline interface

Dutch forensics institute

WhatWorksWhen

StoryTelling

eMetabolomics

Dynaslum

candYgene

eSalsa

EDS

ABCMUSE

Via Appia

3D geospatial

compChemistry

SIMCITY

Point cloud registration

Dirac

KM3NeT

electron scattering

GeoVizualization

Part 2: Radio Astronomy & future infrastructure

Universiteit Leiden
The Netherlands

Radio astronomy science cases

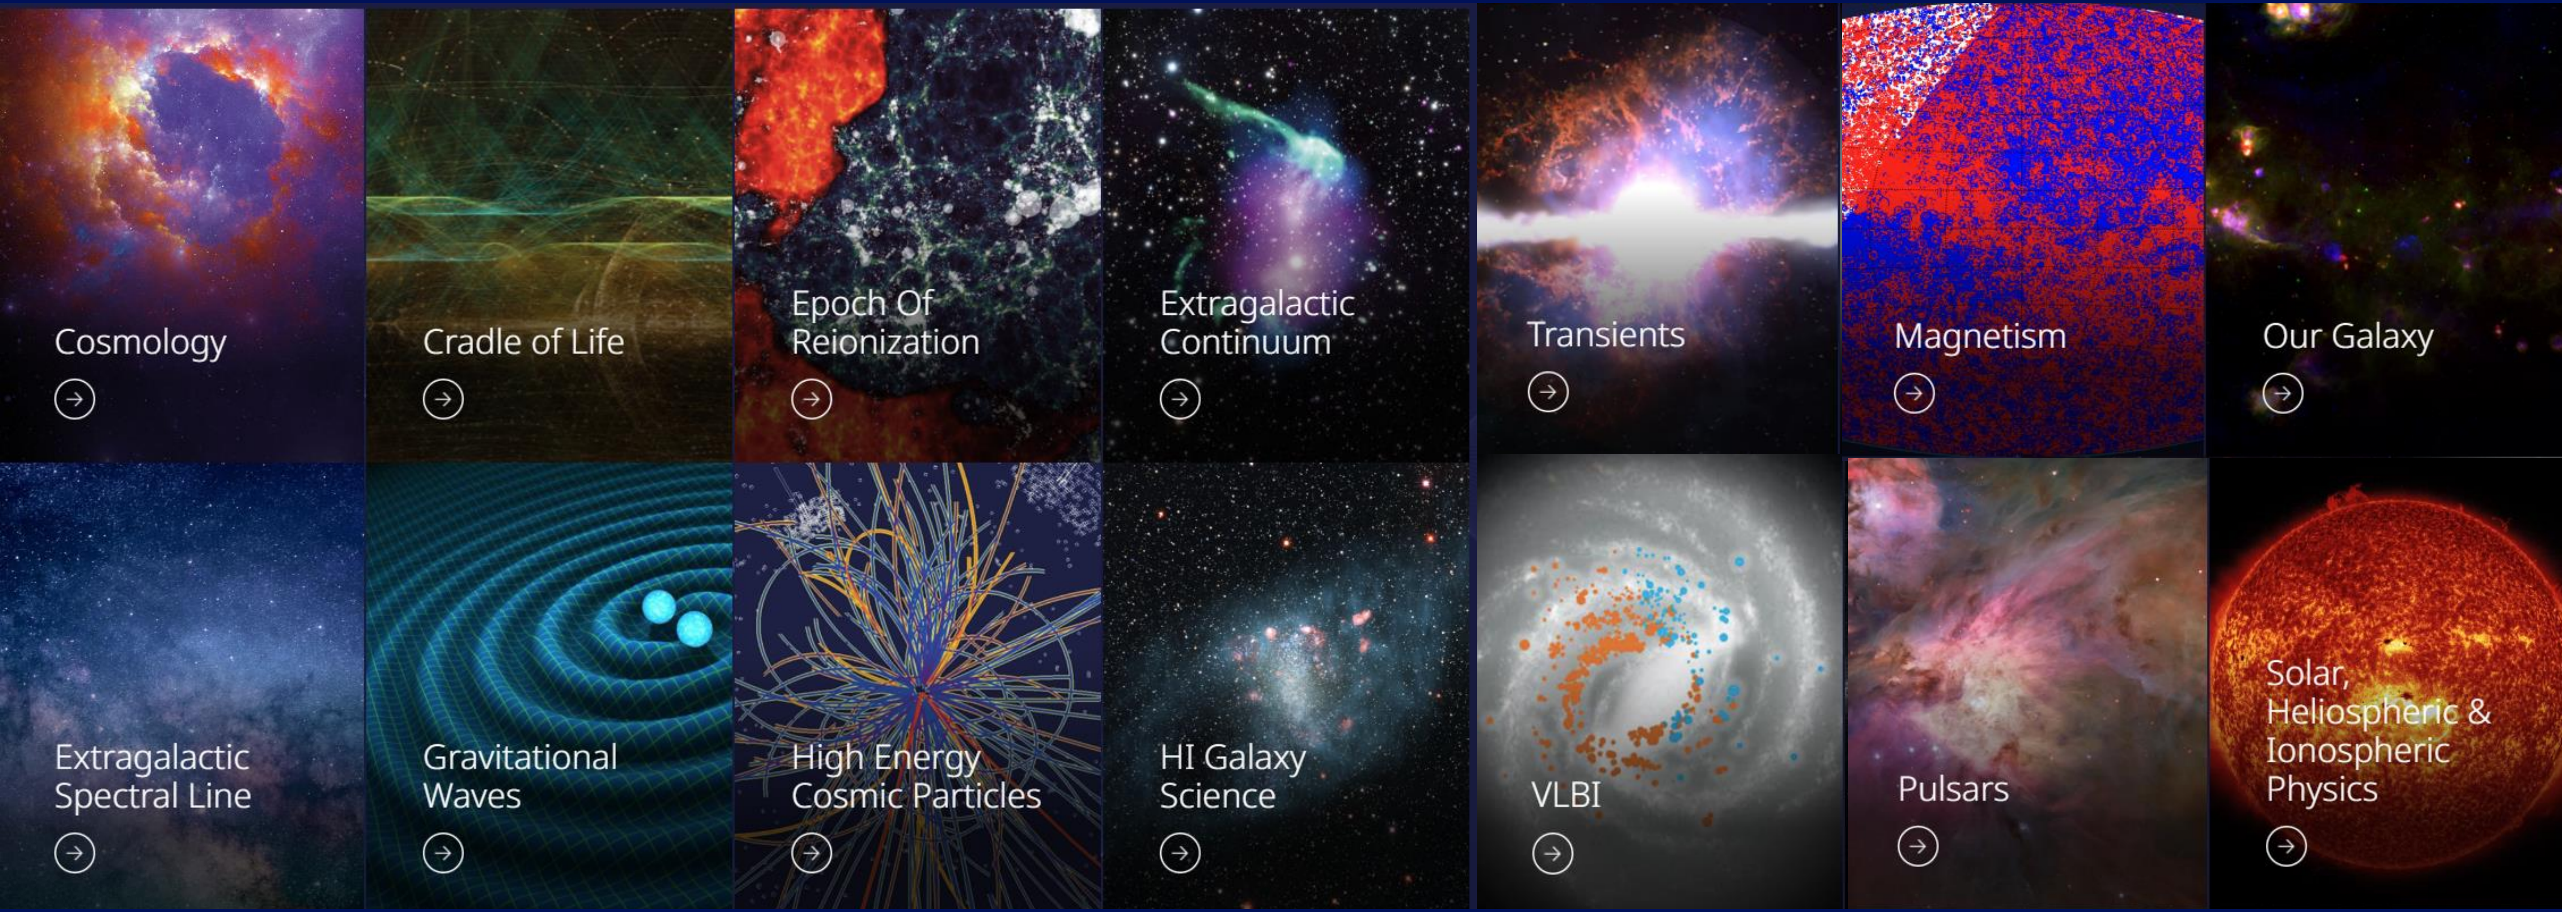

LOFAR

- Largest radio telescope in the world
- ~ 100.000 omni-directional antennas
- 100x more sensitive

- Raw data rate: 10 Tbit/s
- 200 Gbit/s to supercomputer
- petaFLOPS

Hercules A. Credit: R. Timmerman; LOFAR & Hubble Space Telescope

LOFAR low-band antennas

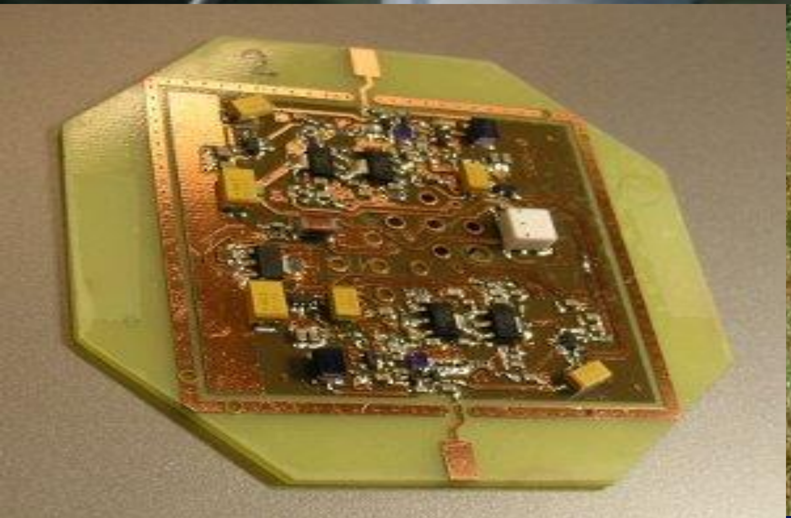

LOFAR high-band antennas

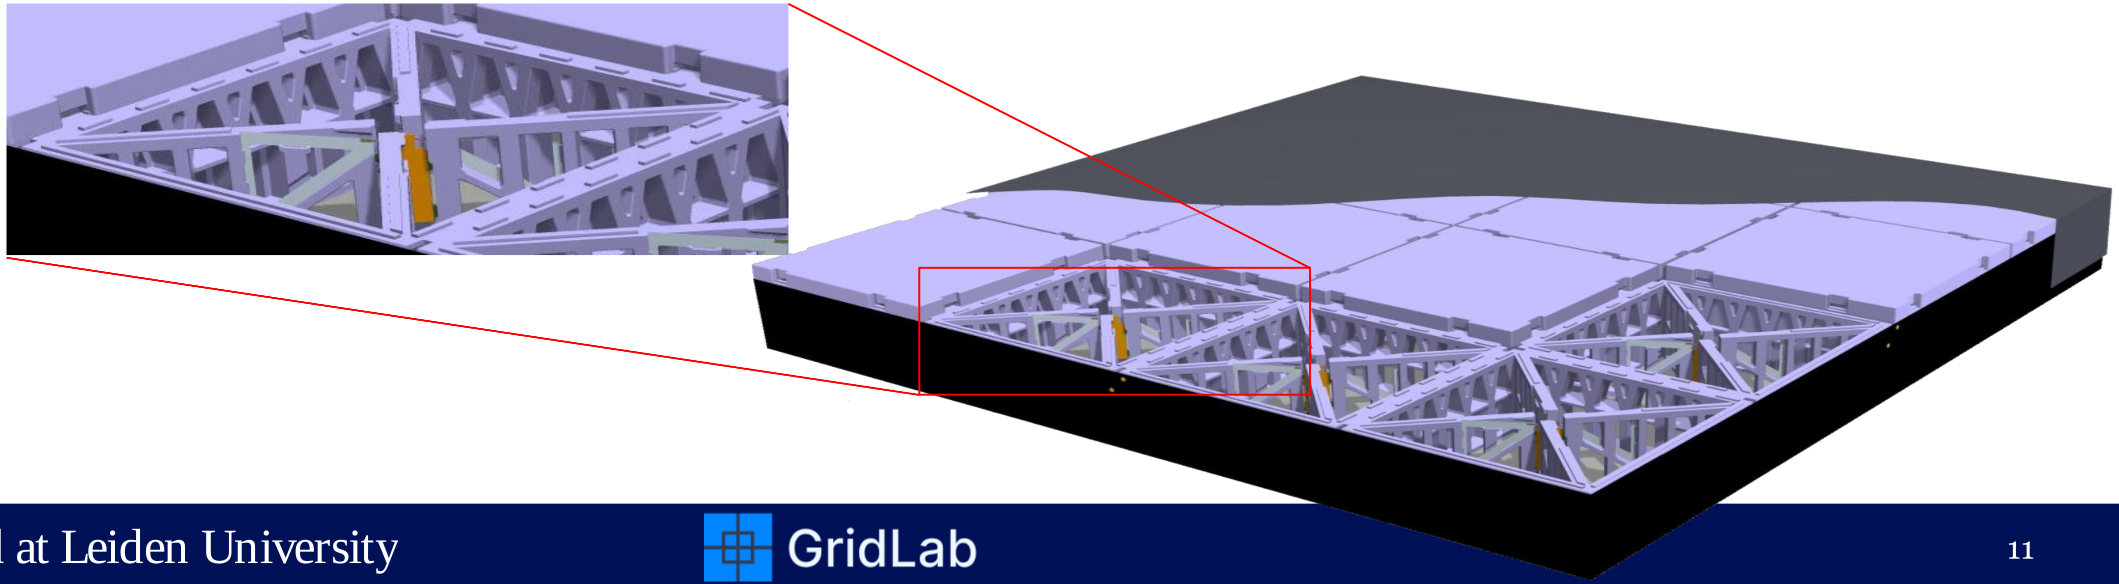

LOFAR Station (150m)


~ 3 km


Onsala
Irbene
Birr
Dutch stations
Chilbolton
Norderstedt
Bałdy
Potsdam
Borówiec
Jülich
Effelsberg
Tautenburg
Łazy
Nançay
Unterweilenbach
Medicina

LOFAR pipelines

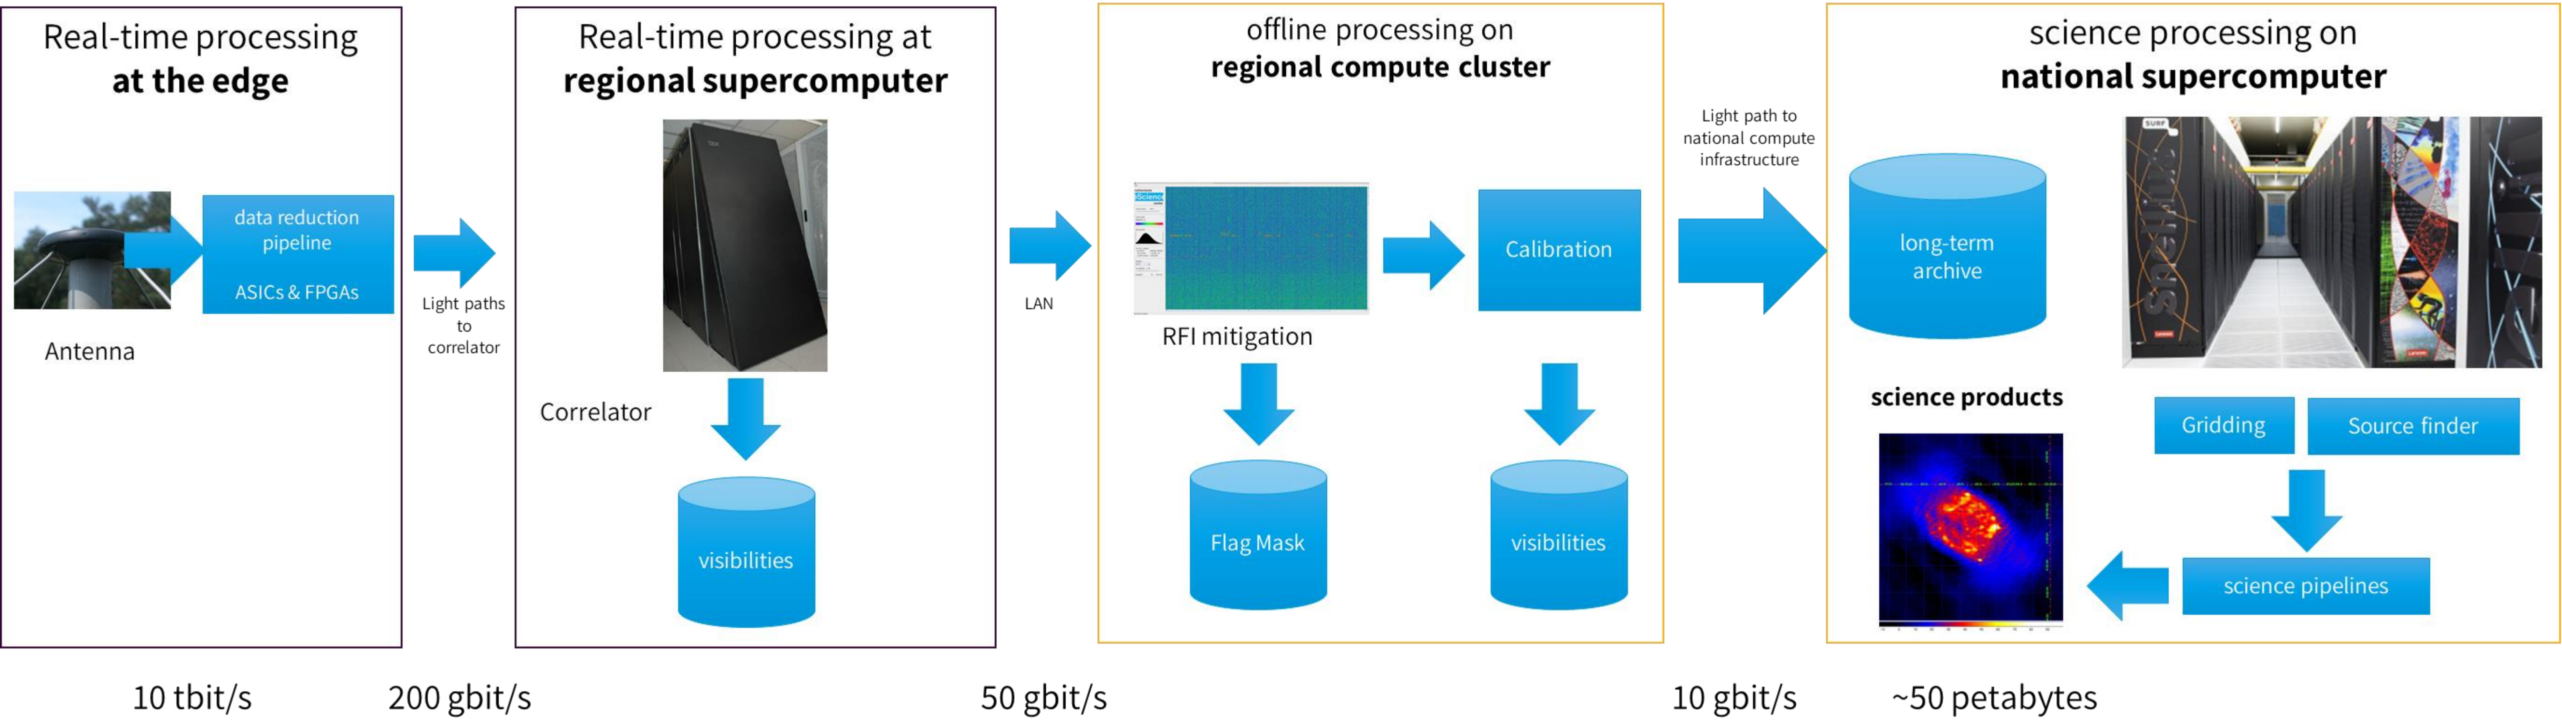

The next generation: the SKA

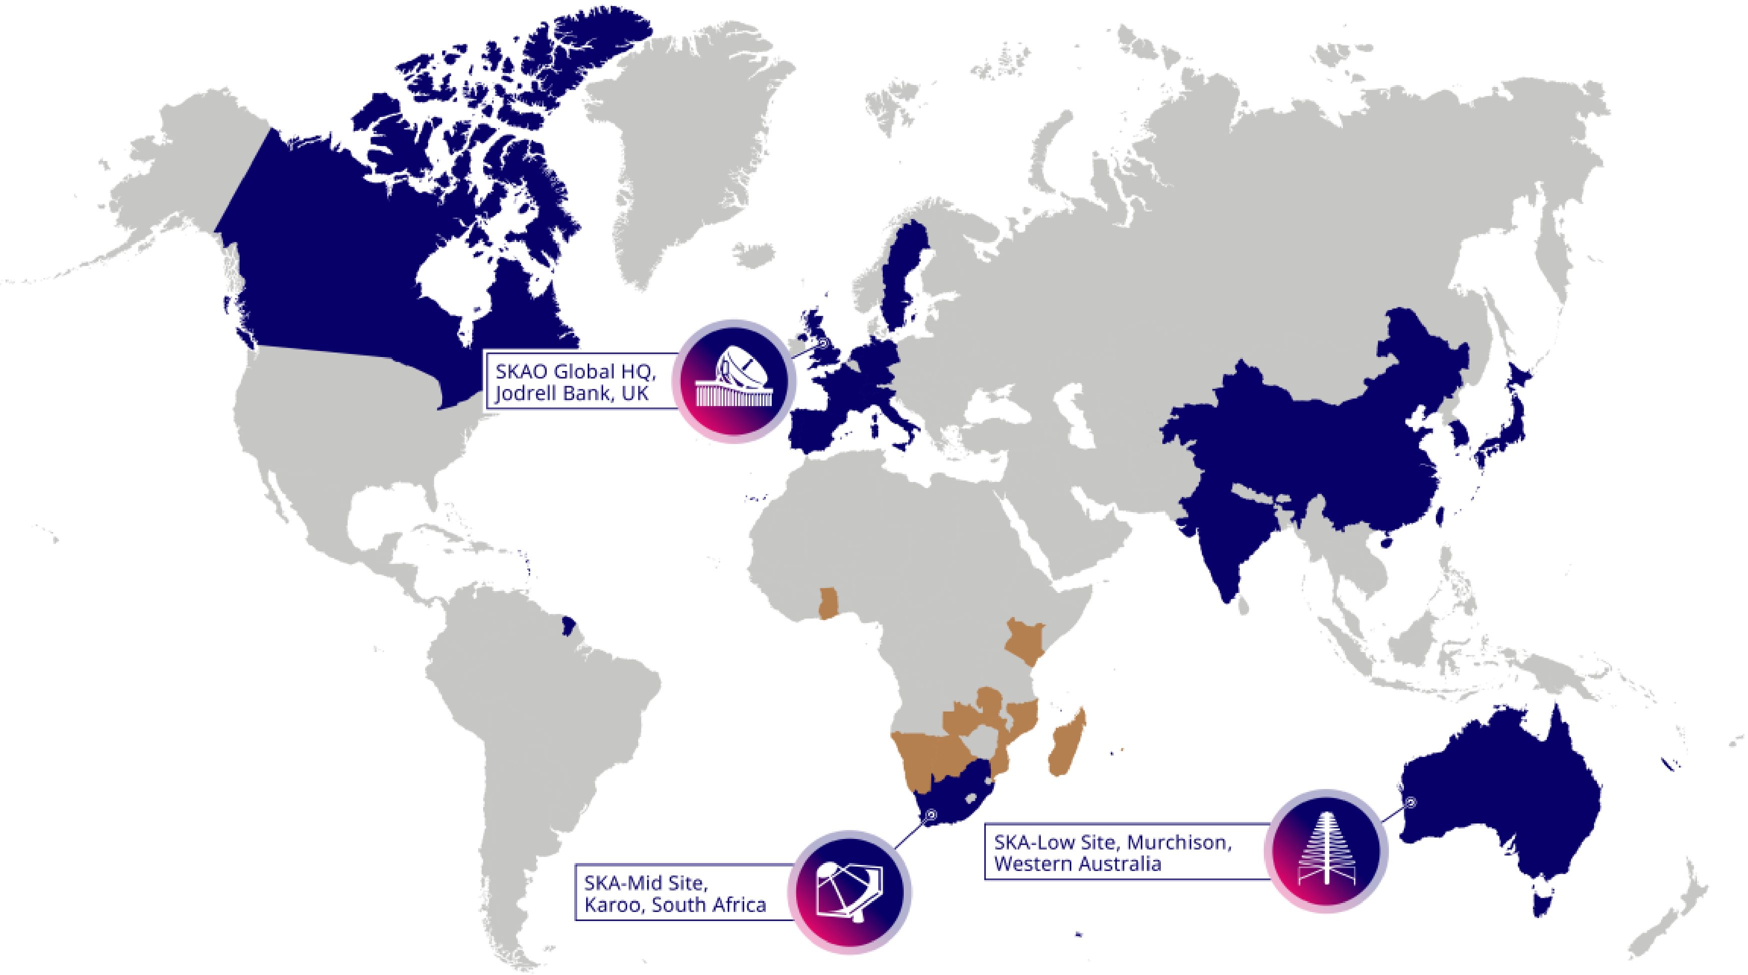

SKAO Partnership - includes SKAO Member States* and SKAO Observers (as of July 2023)

African Partner Countries

SKA-mid

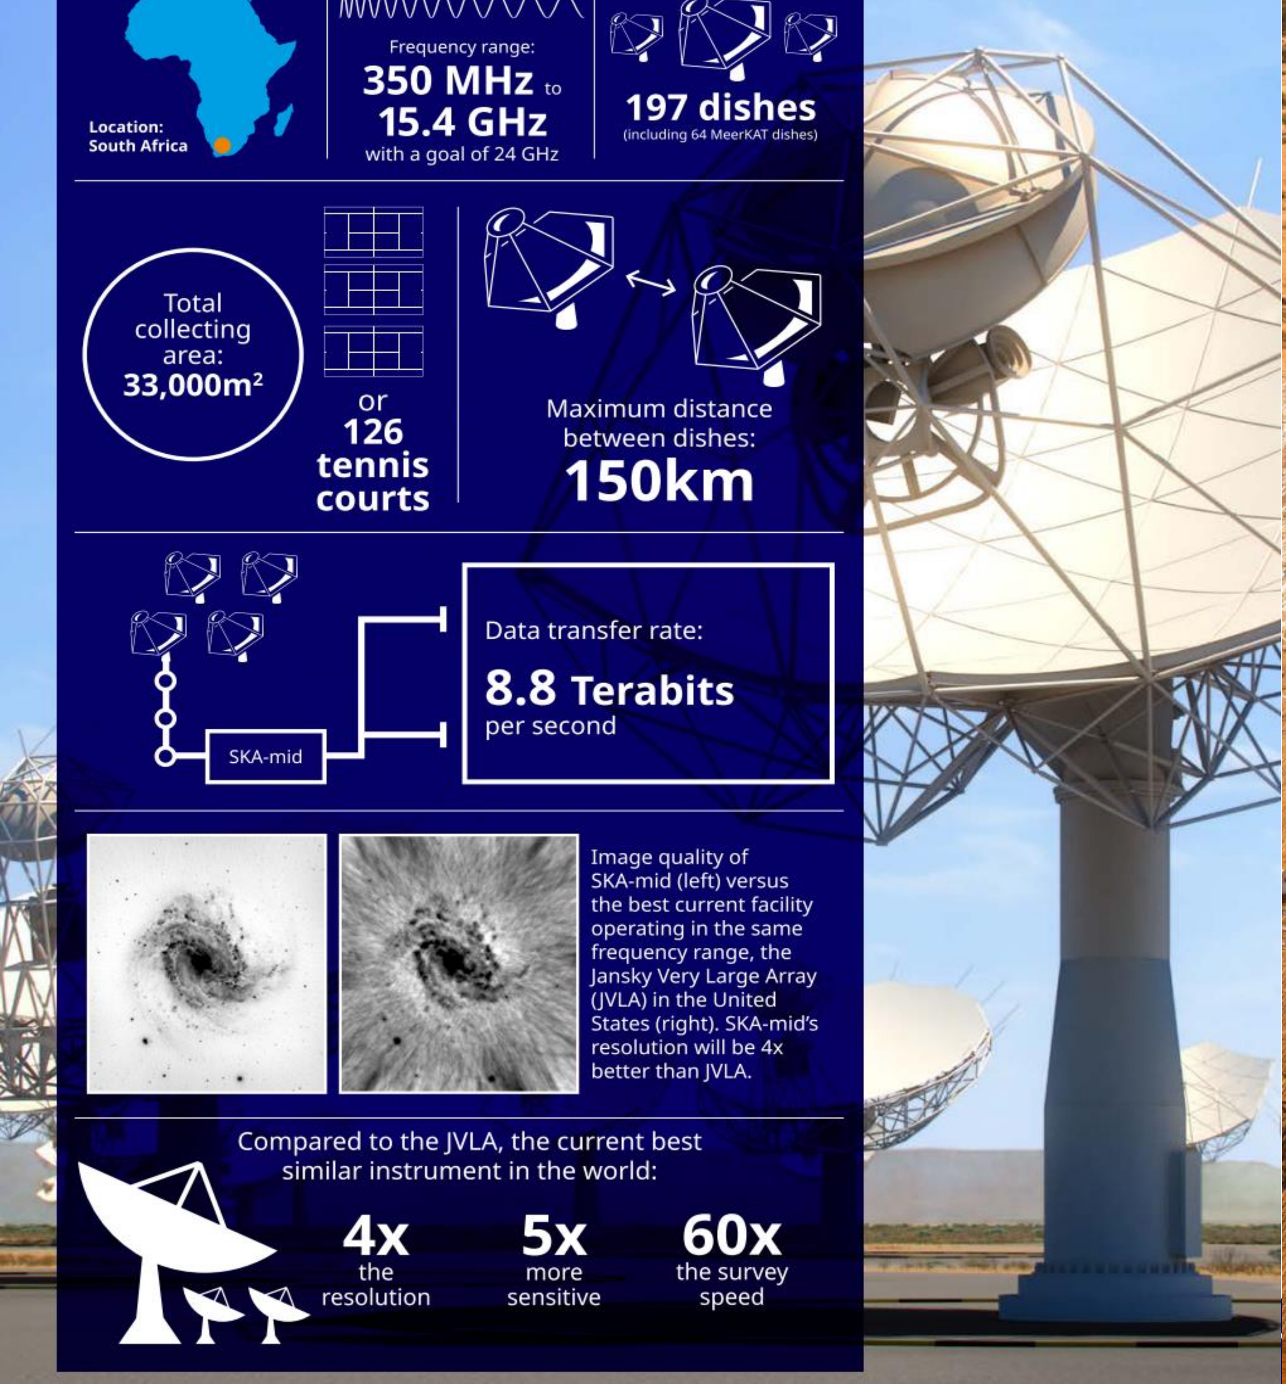

SKA-low

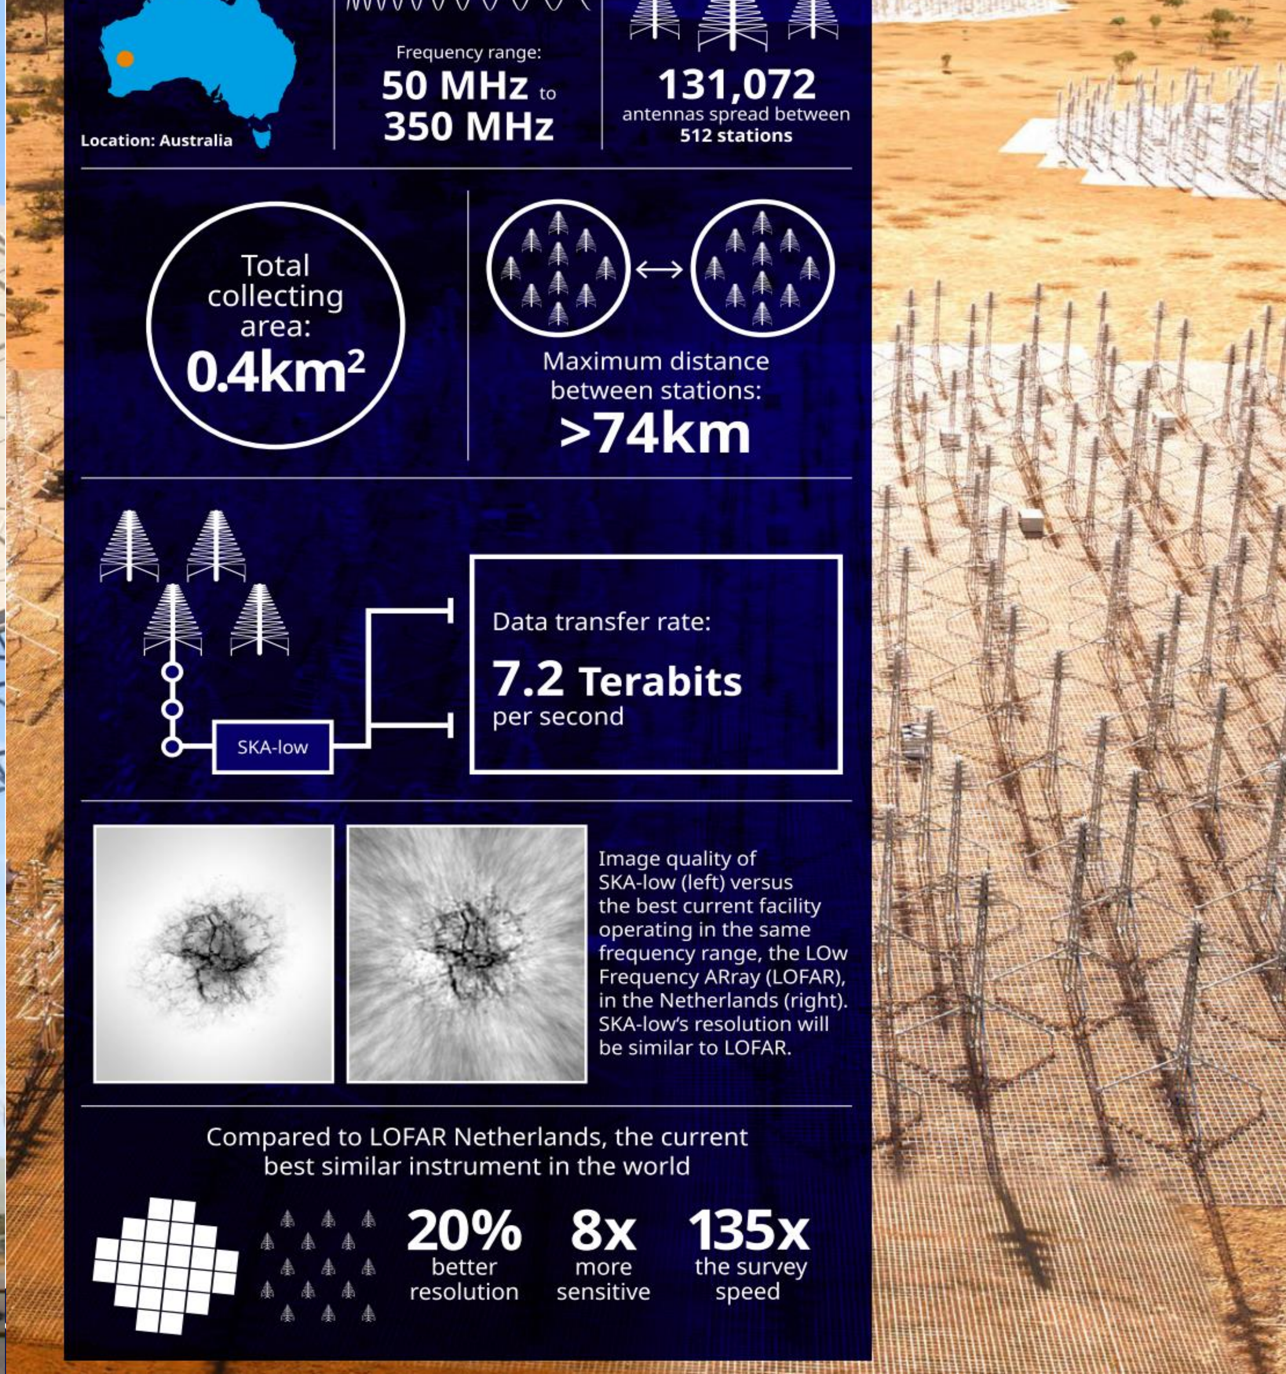

SKA: A distributed system at the exascale

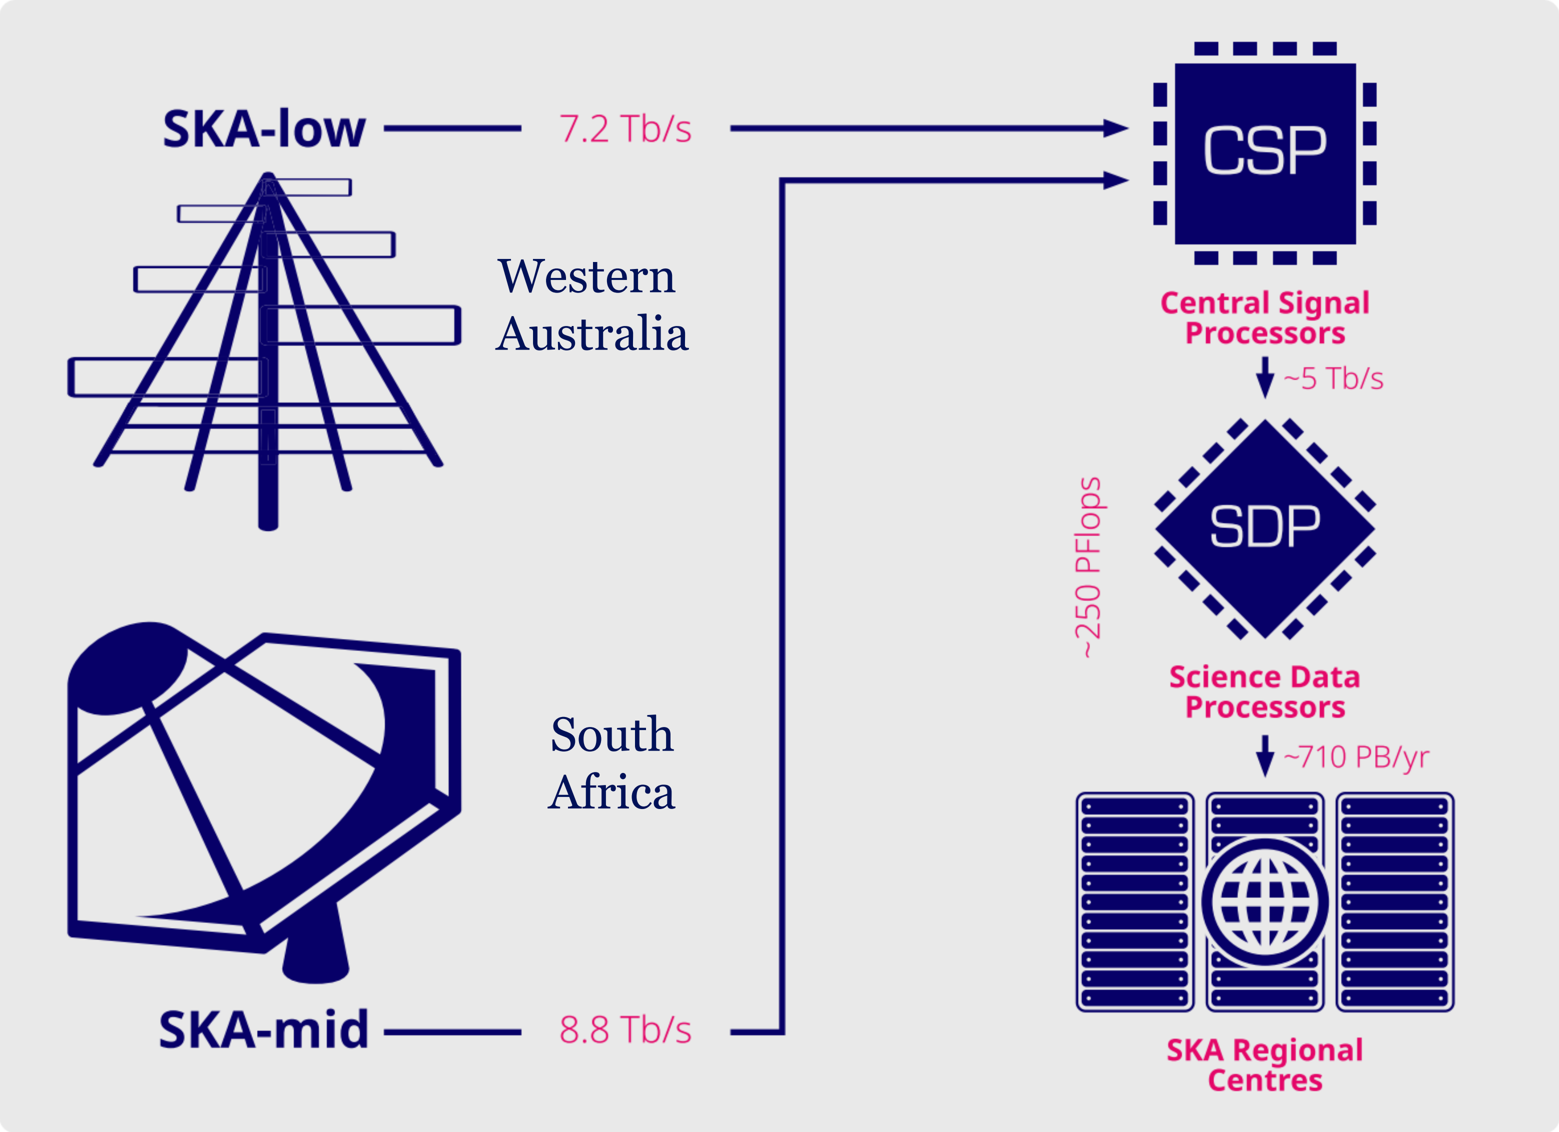

A typical image is 32k x 32k x 64k

Archive:
Many exabytes in the lifetime of the instrument

Oh dear!

This is only phase 1.

Phase 2 is envisioned to be 10-20x bigger!

… And some algorithms are quadratic…

What does this require from the future infrastructure?

- Vast amounts of data transport
- Data archival: metadata and user-friendly search interface / visualization

- More processing with less energy in the entire continuum
  - Edge processing at the receivers. In the middle of nowhere → power budget very small
  - CSP: more power budget, but huge data volume, processing is **quadratic** in the number of receivers
  - SDP: still huge data volume, but much more complex processing (calibration, imaging)
  - Regional centers: very diverse and complex processing workflows, with very different hardware requirements

- Flexibility
  - Vastly different science cases
  - Telescopes have a lifetime measured in decades. Unforeseen science cases **will happen**
  - Algorithms also improve continuously. Example: novel calibration with AI-based methods → completely different requirements

- Robustness
  - This is an operational and real-time system, needs to **work** 24/7
  - Infrastructure needs to be maintained and needs to evolve seamlessly for decades, also to adapt to the science cases

Future technologies that may help (or get in the way)

Hardware

- Specialized accelerators
- Neuromorphic
- Quantum radio astronomy

Virtualized & programmable everything

- Router, switch, NIC
- Storage layers, SSD
- Cache hierarchy
- Enabling:
  - Initial processing in the network
  - Streaming data from instrument directly into accelerator without any host involvement
  - Energy savings

Logging, monitoring & measuring everything

- Better system health management, robust applications
- More grip on energy usage

Software

- More GPU codes
  - Currently polyphase filter, correlator, beam former, imaging
  - Many science pipelines not yet re-implemented
- Code generation with large language models
- More AI-based algorithms
  - Currently calibration, RFI mitigation, some science cases

Ease of use may be more important than an optimal solution

Structural funding for software infrastructure

A concrete GridLab++ use case?

Problem

- System Health Management
  - Hundreds of thousands of components
  - Spanning the entire compute continuum
- Failures (harsh conditions)
  - antennas
  - radio frequency interference (RFI)
  - all processing elements (edge, local super, cluster, national super, …)
  - memory
  - storage
  - networks
  - software
  - …

Solution

- Detect anomalies / failures
- Find root cause
- Mitigate

Challenges

- combine classical methods & AI
- throughout the continuum (distributed system)
- real-time system
- limited resources
- large data volumes
- heterogeneous data
  - data itself, monitoring info, system logs, software
- integration and deployment

Questions ?

Universiteit Leiden
The Netherlands

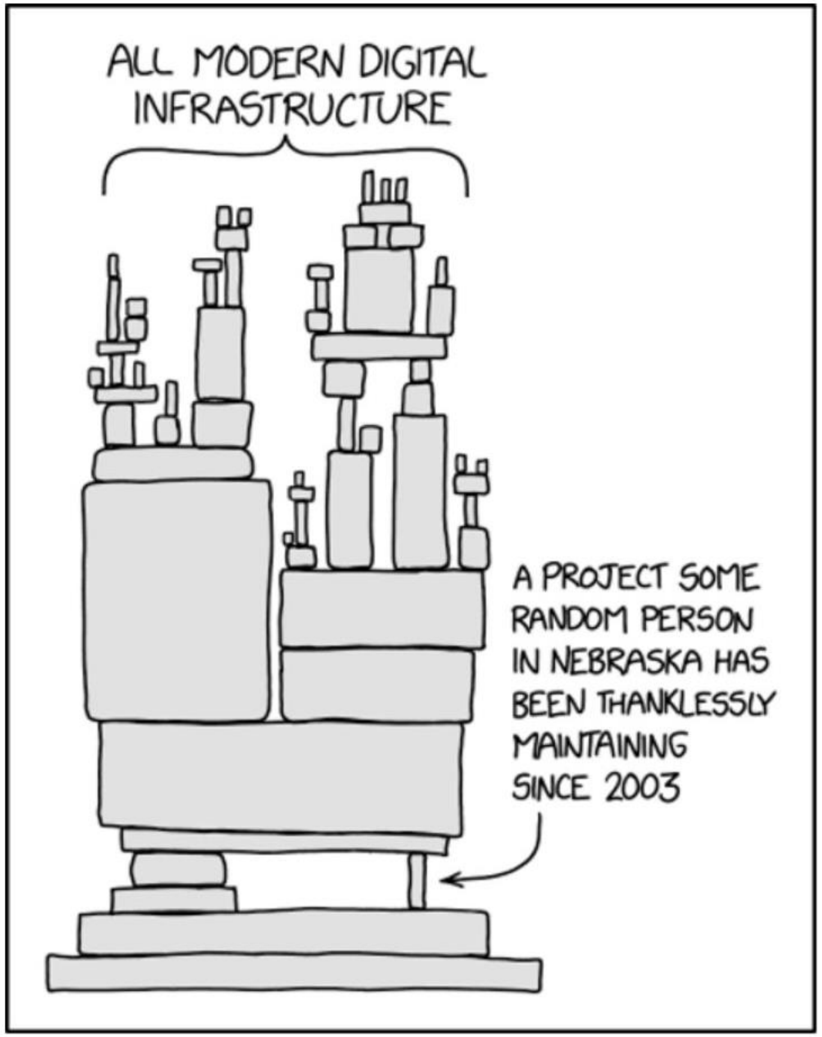

Randall Munroe, *Dependency*, xkcd, https://xkcd.com/2347/ (last visited Aug. 18, 2021) (licensed under CC BY-NC 2.5).

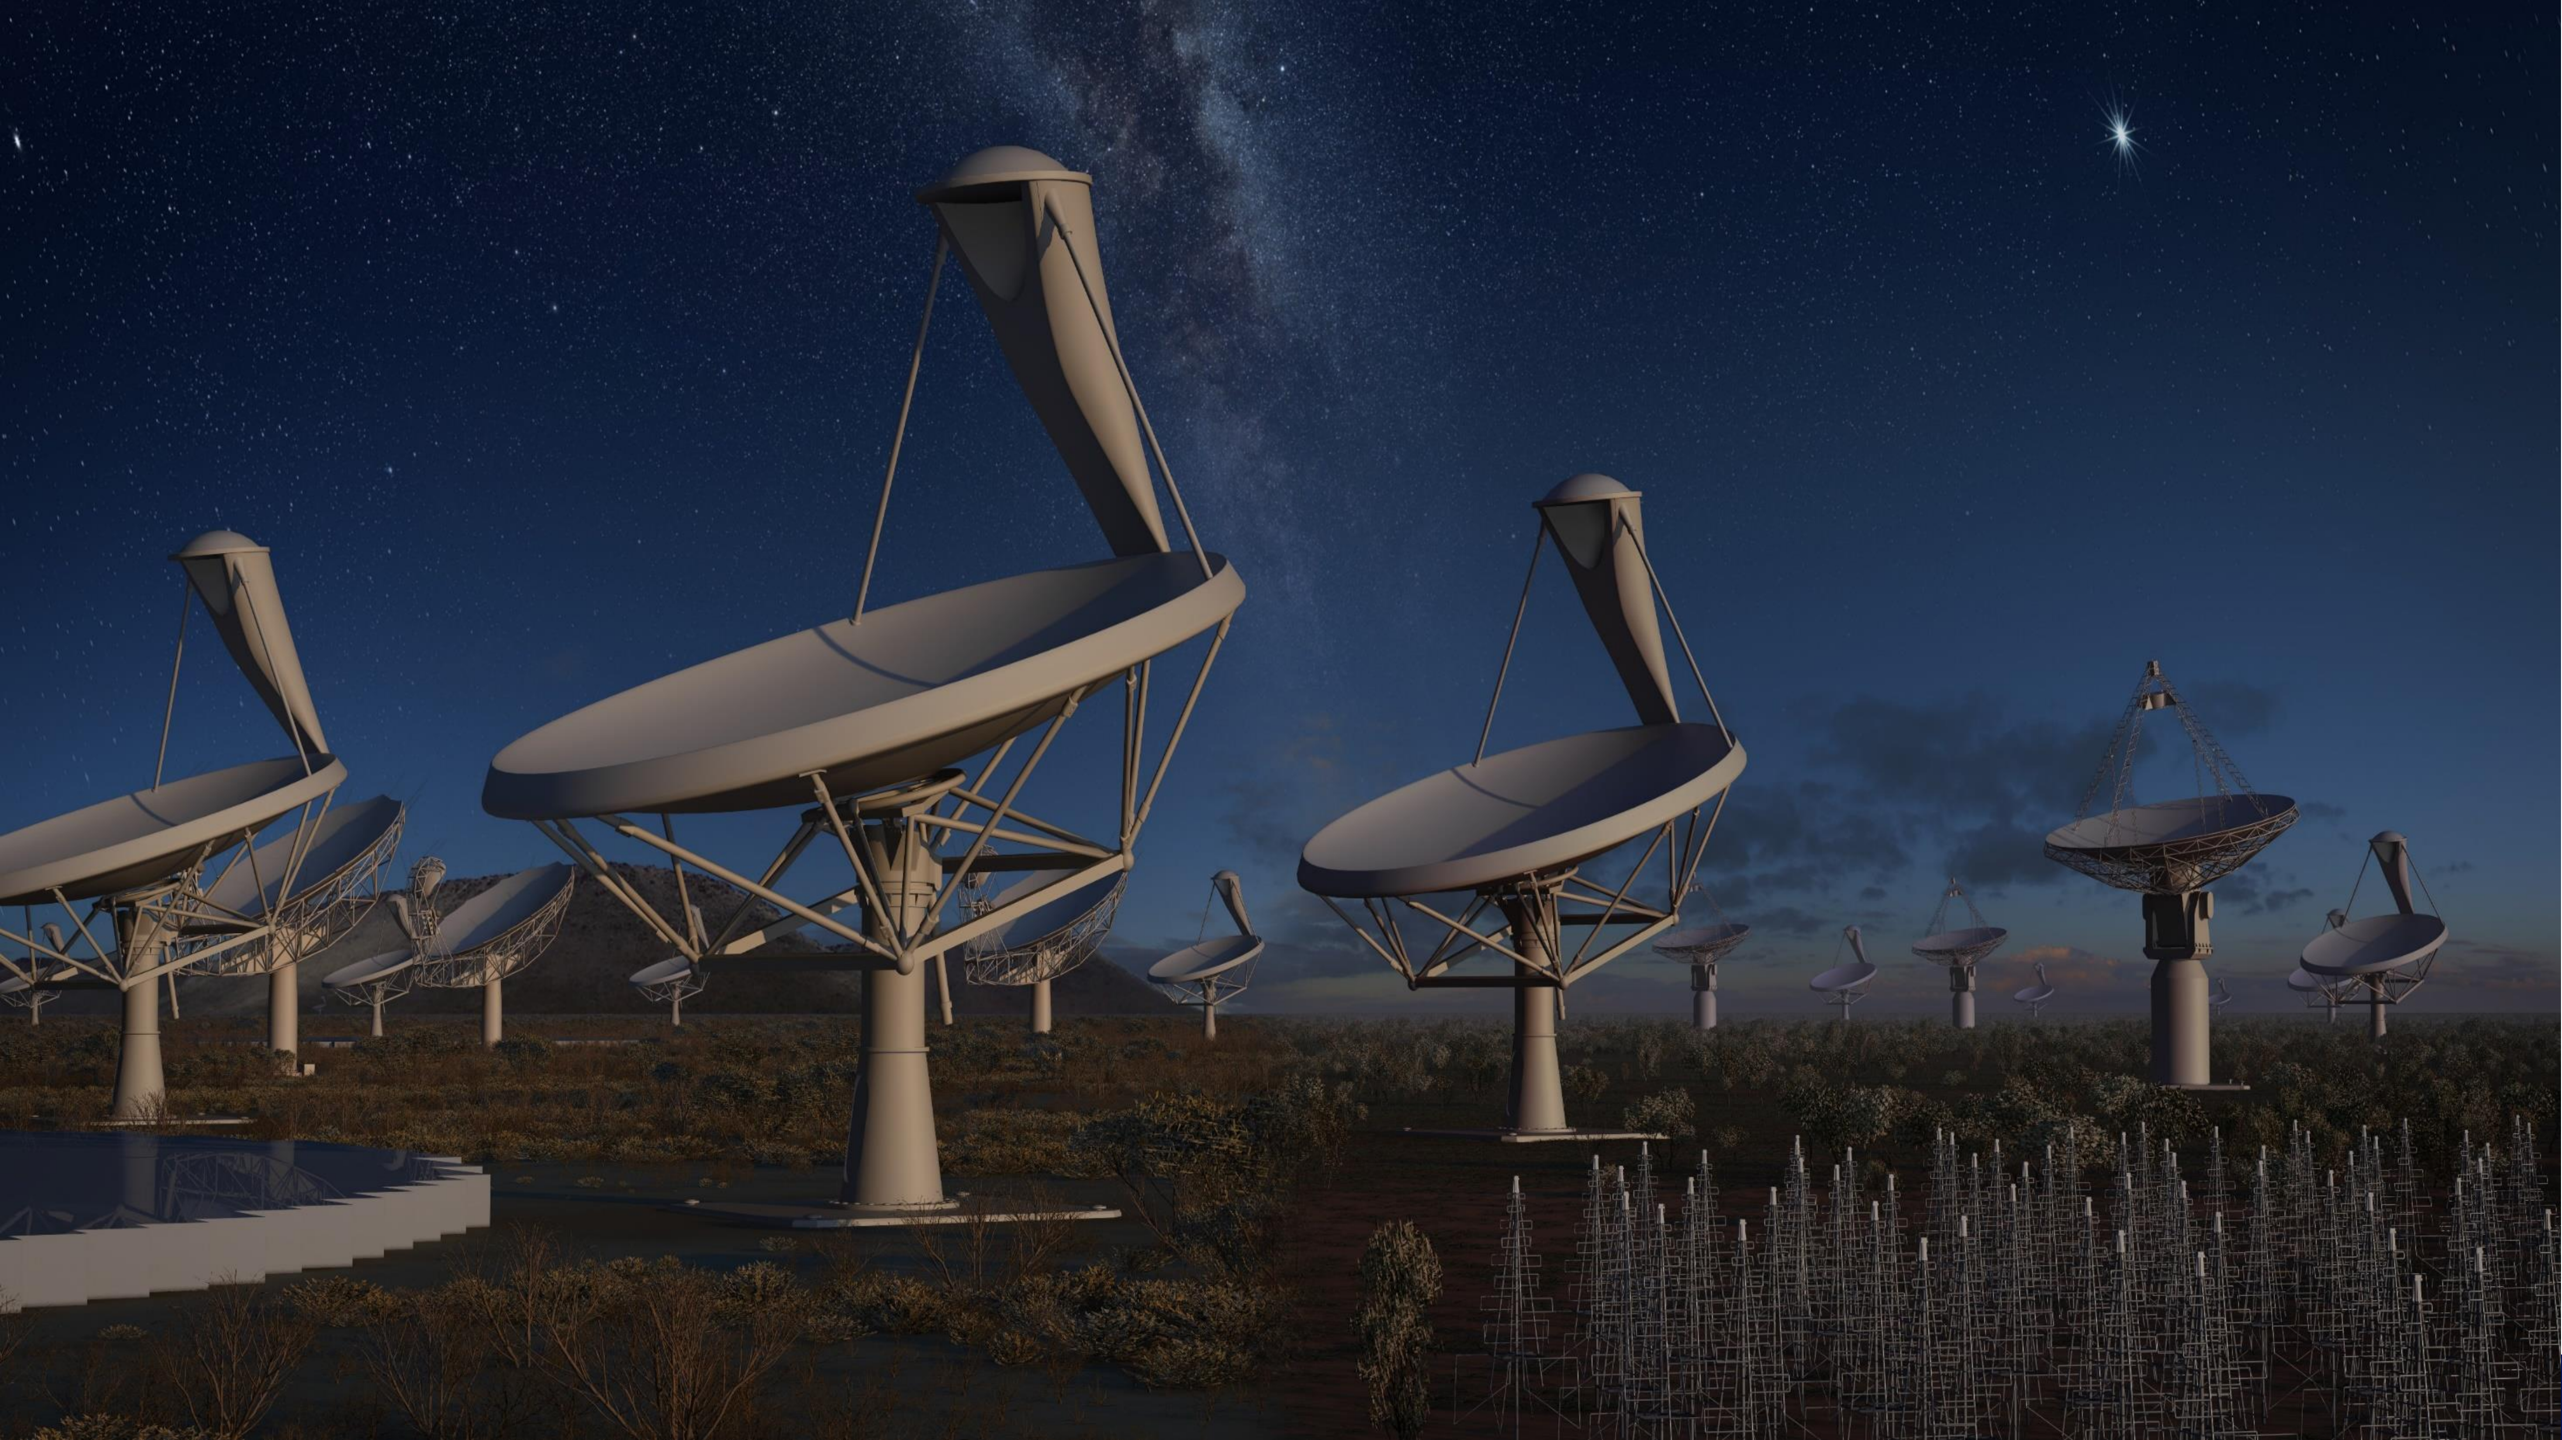

Three LOFAR pipelines

Transient pipeline:
imaging in real-time

Antenna

Light paths to
correlator

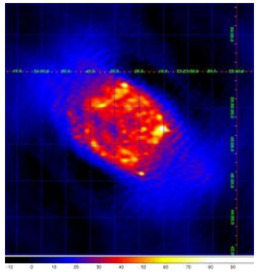

Imaging mode

Source finder

catalog

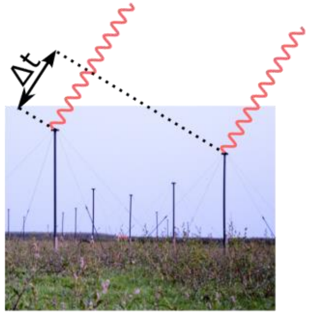

Beam forming mode

Pulsar pipeline

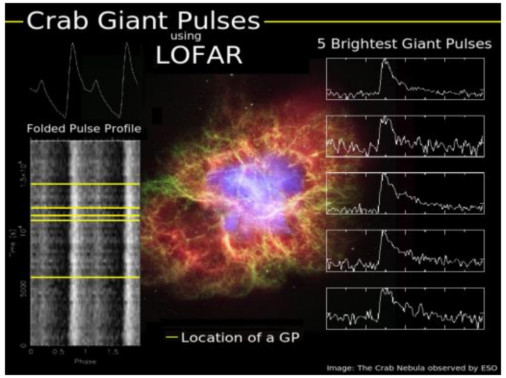

RFI mitigation

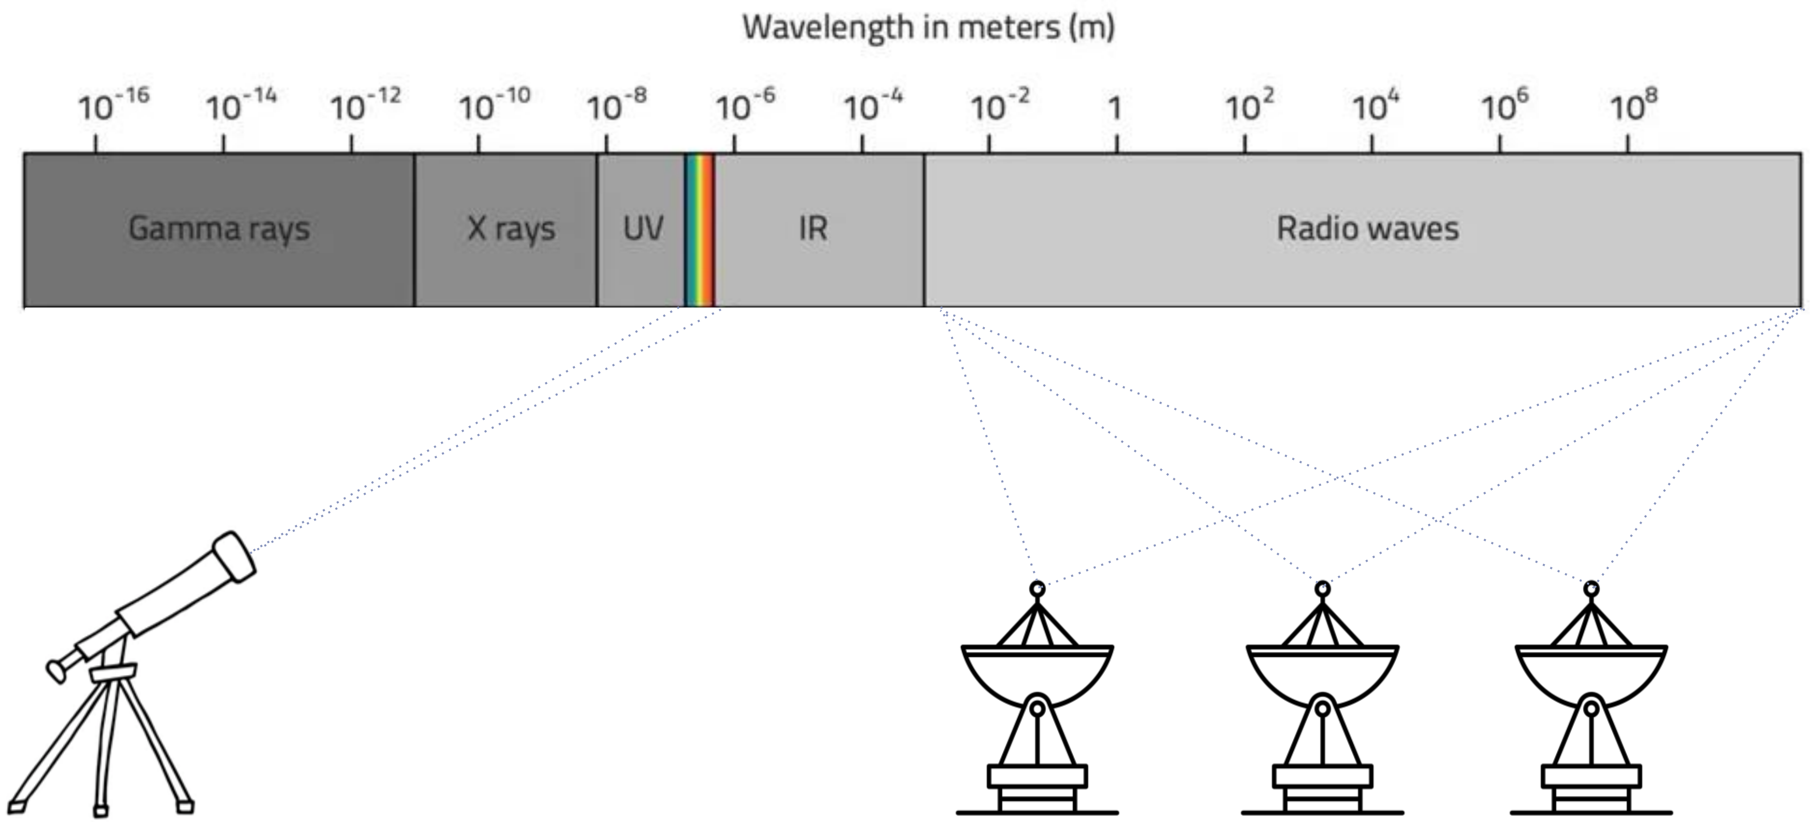

RFI mitigation

Figure 1 - Messier 83 in visible spectrum

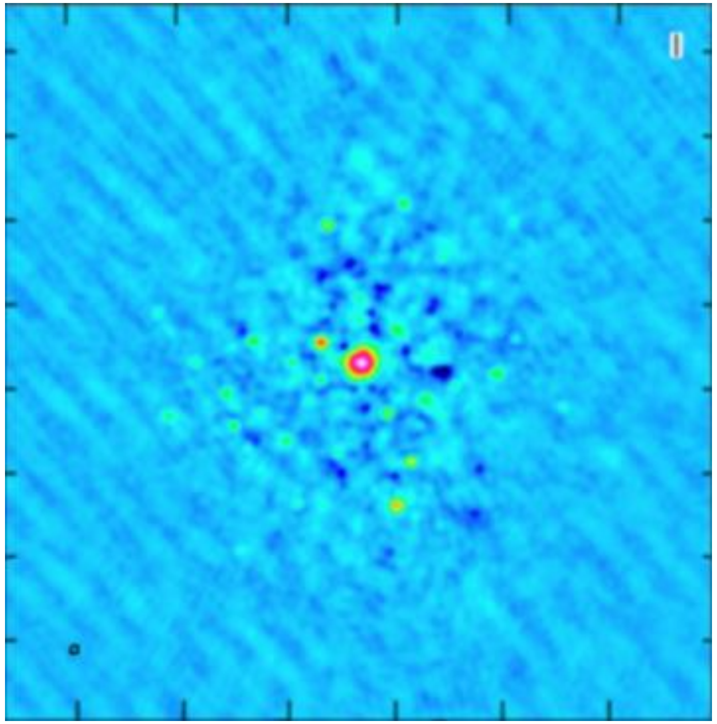

Figure 2 - Messier 83 in radio spectrum contaminated with RFI (after imaging)

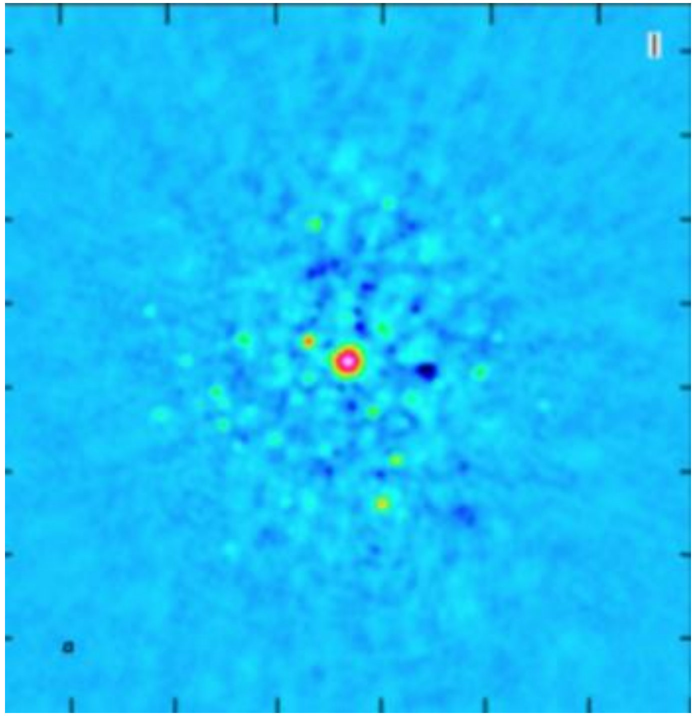

Figure 3 - Messier 83 in radio spectrum with RFI blanked (after imaging)