An Update on MeerKAT*

Bradley Frank
Russ Taylor, Rob Simmonds, Fernando Camilo

*& IDIA

IDIA Inter-University Institute for Data Intensive Astronomy
UNIVERSITY OF CAPE TOWN
IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD
ARC AFRICAN RESEARCH CLOUD
SKA AFRICA SQUARE KILOMETRE ARRAY

Overview

MeerKAT

MeerKAT Survey(s)

Science with MeerKAT Surveys

MeerKAT

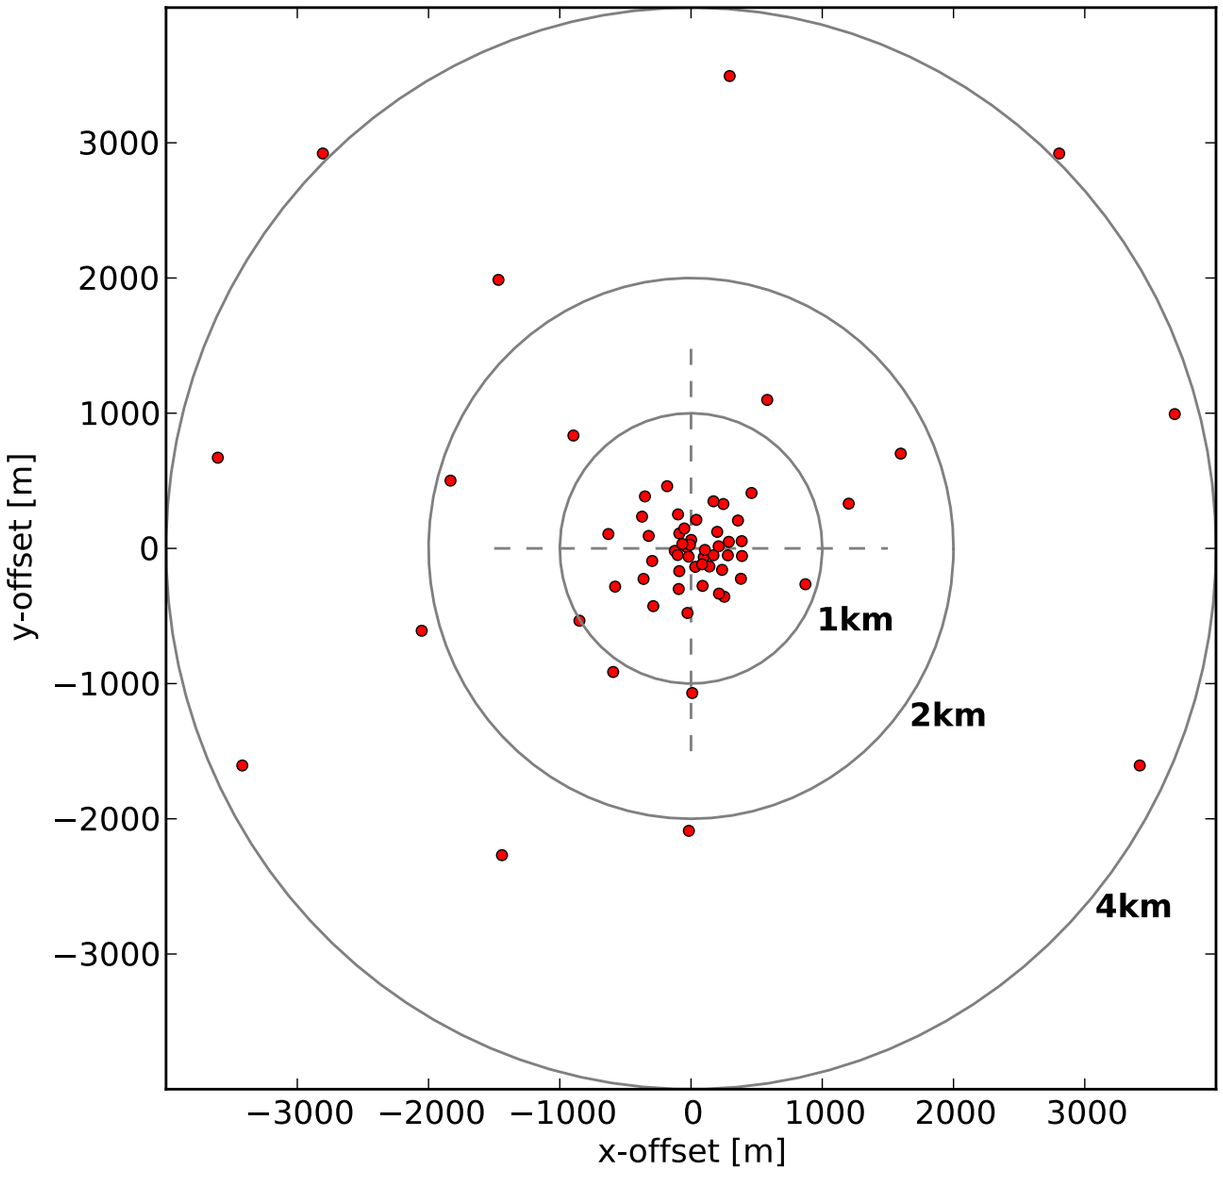

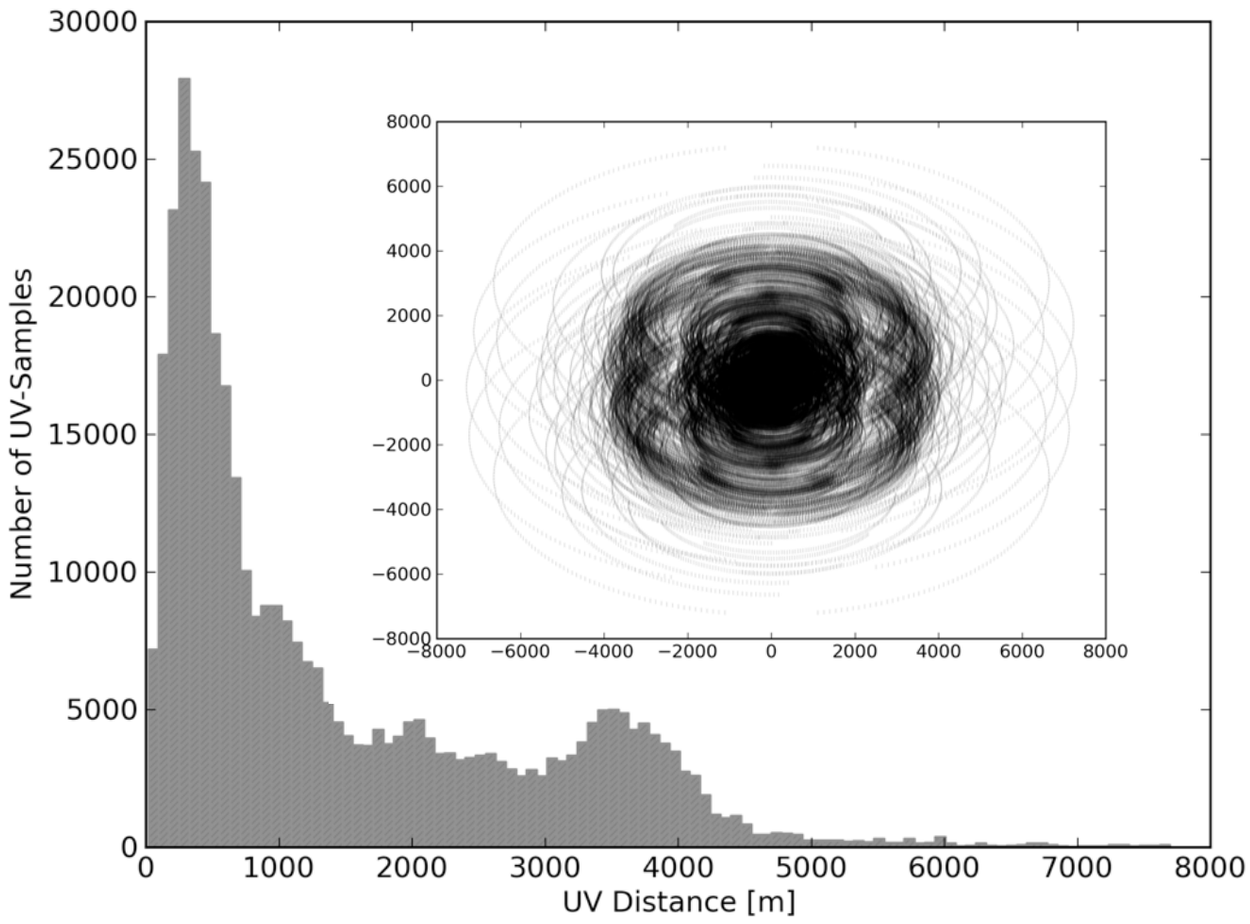

- 64x 13.5m Offset Gregorian Dishes.
  - L-band (0.9-1.67GHz)
  - UHF (0.58-1.0GHz)
  - S-band (1.75-3.5GHz)

MeerKAT

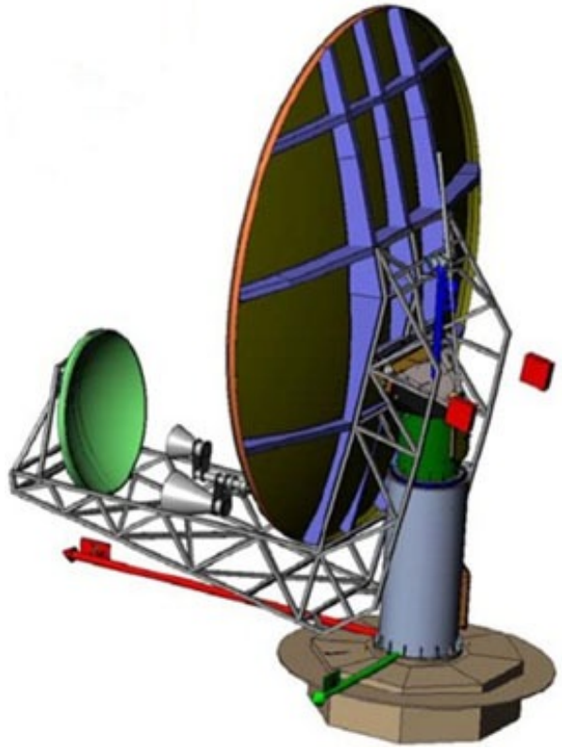

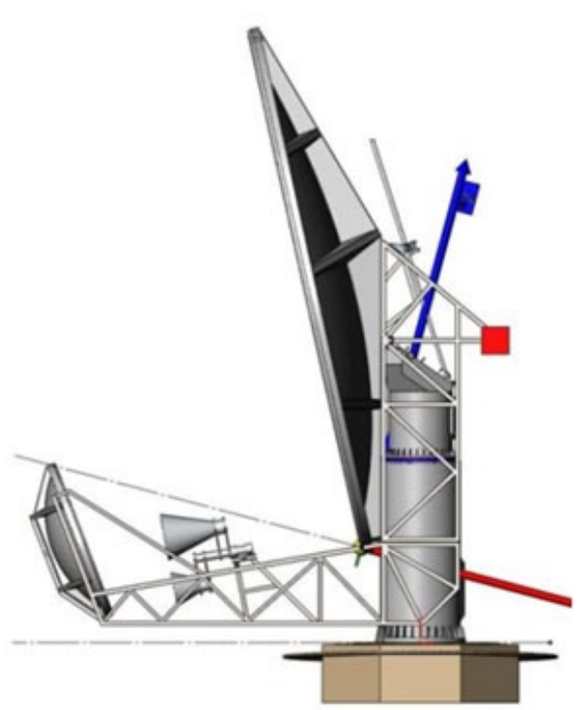

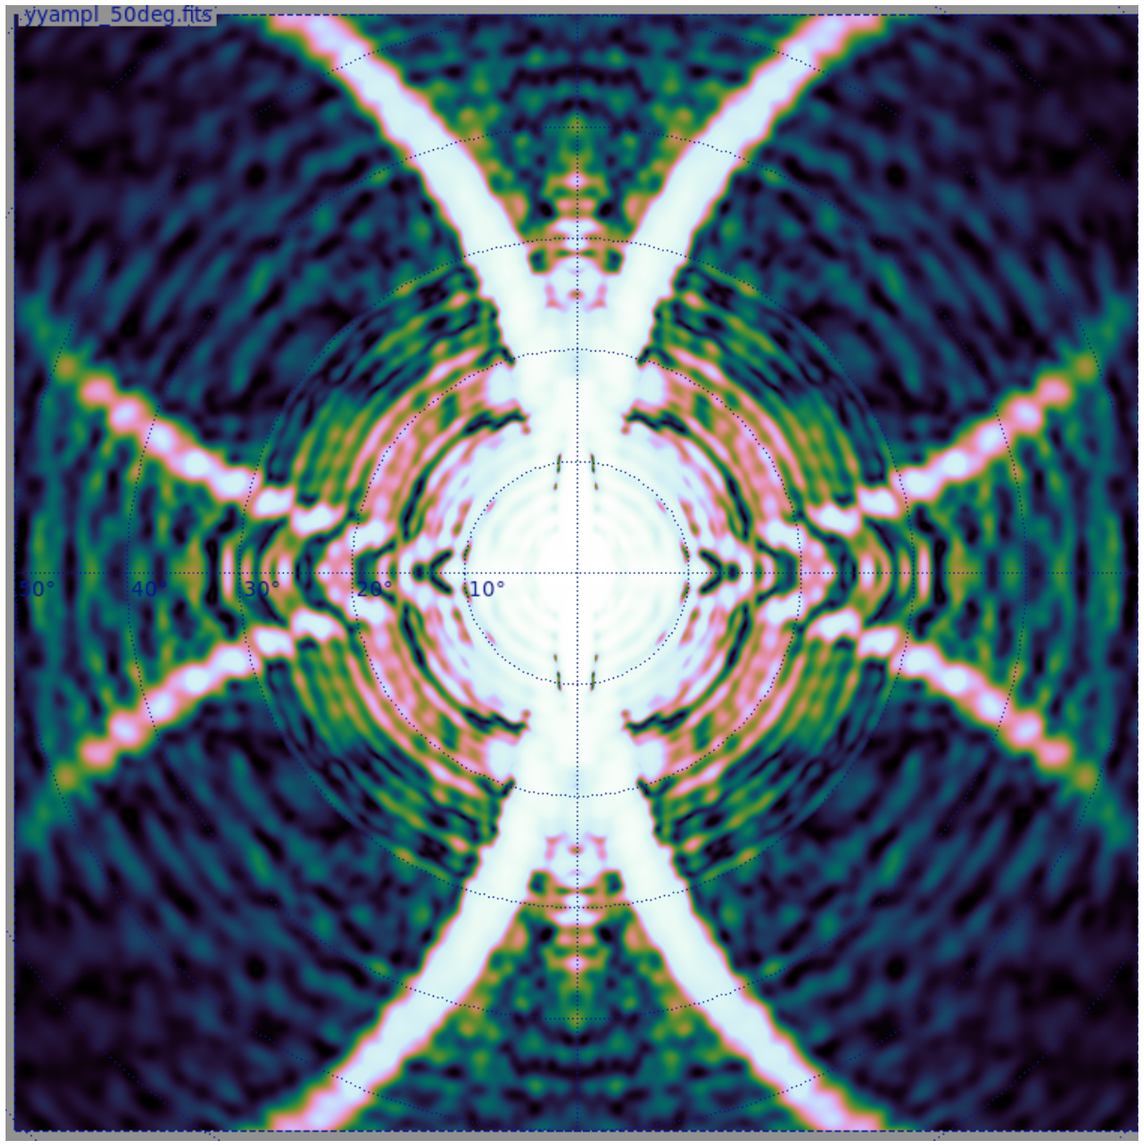
yyampl_50deg.fits
50°
40°
30°
20°
10°

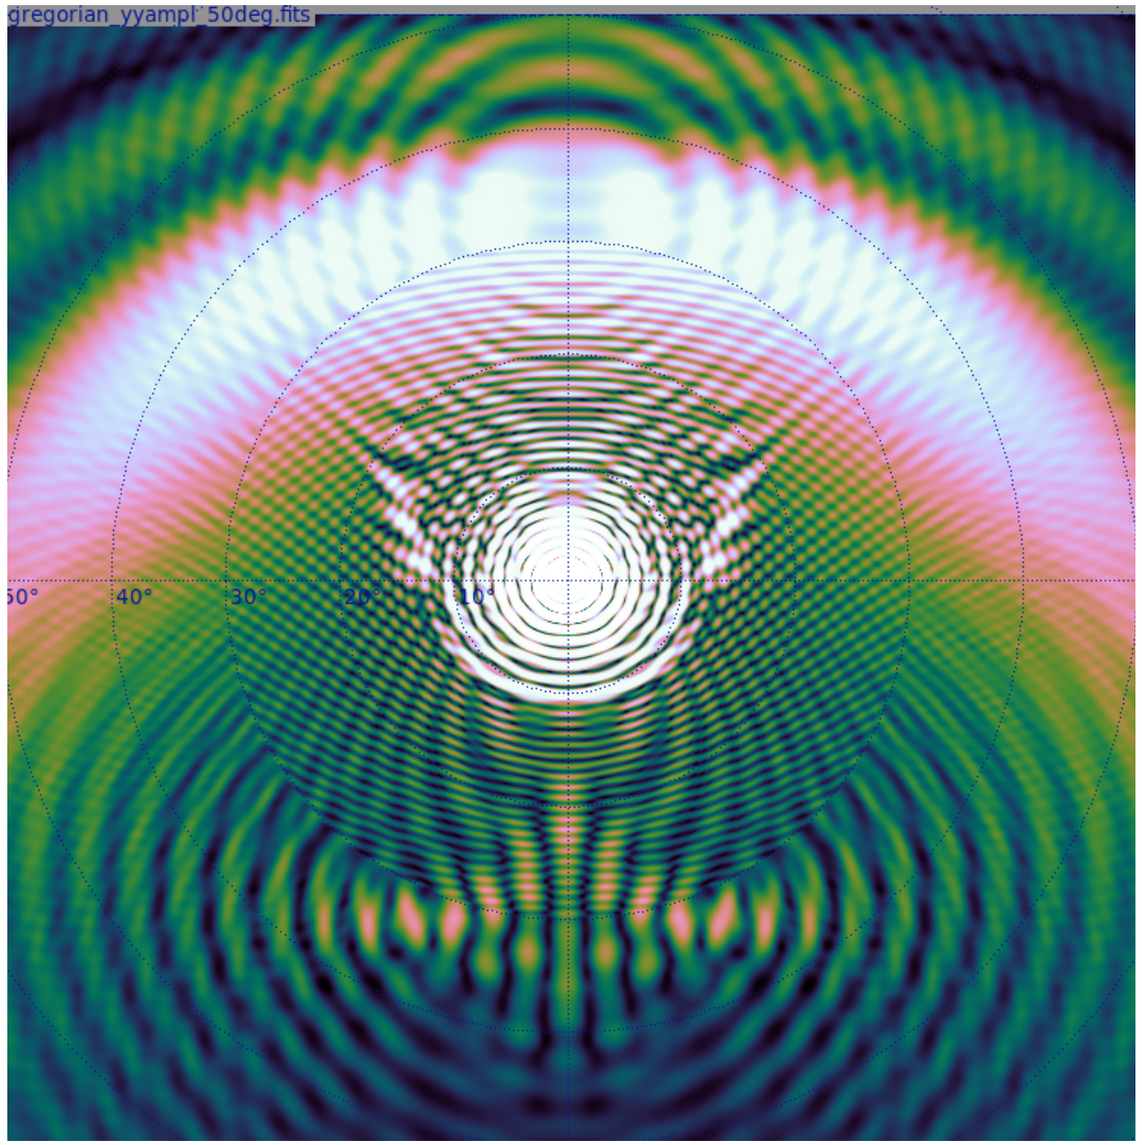
gregorian_yyampl_50deg.fits
50°
40°
30°
20°
10°

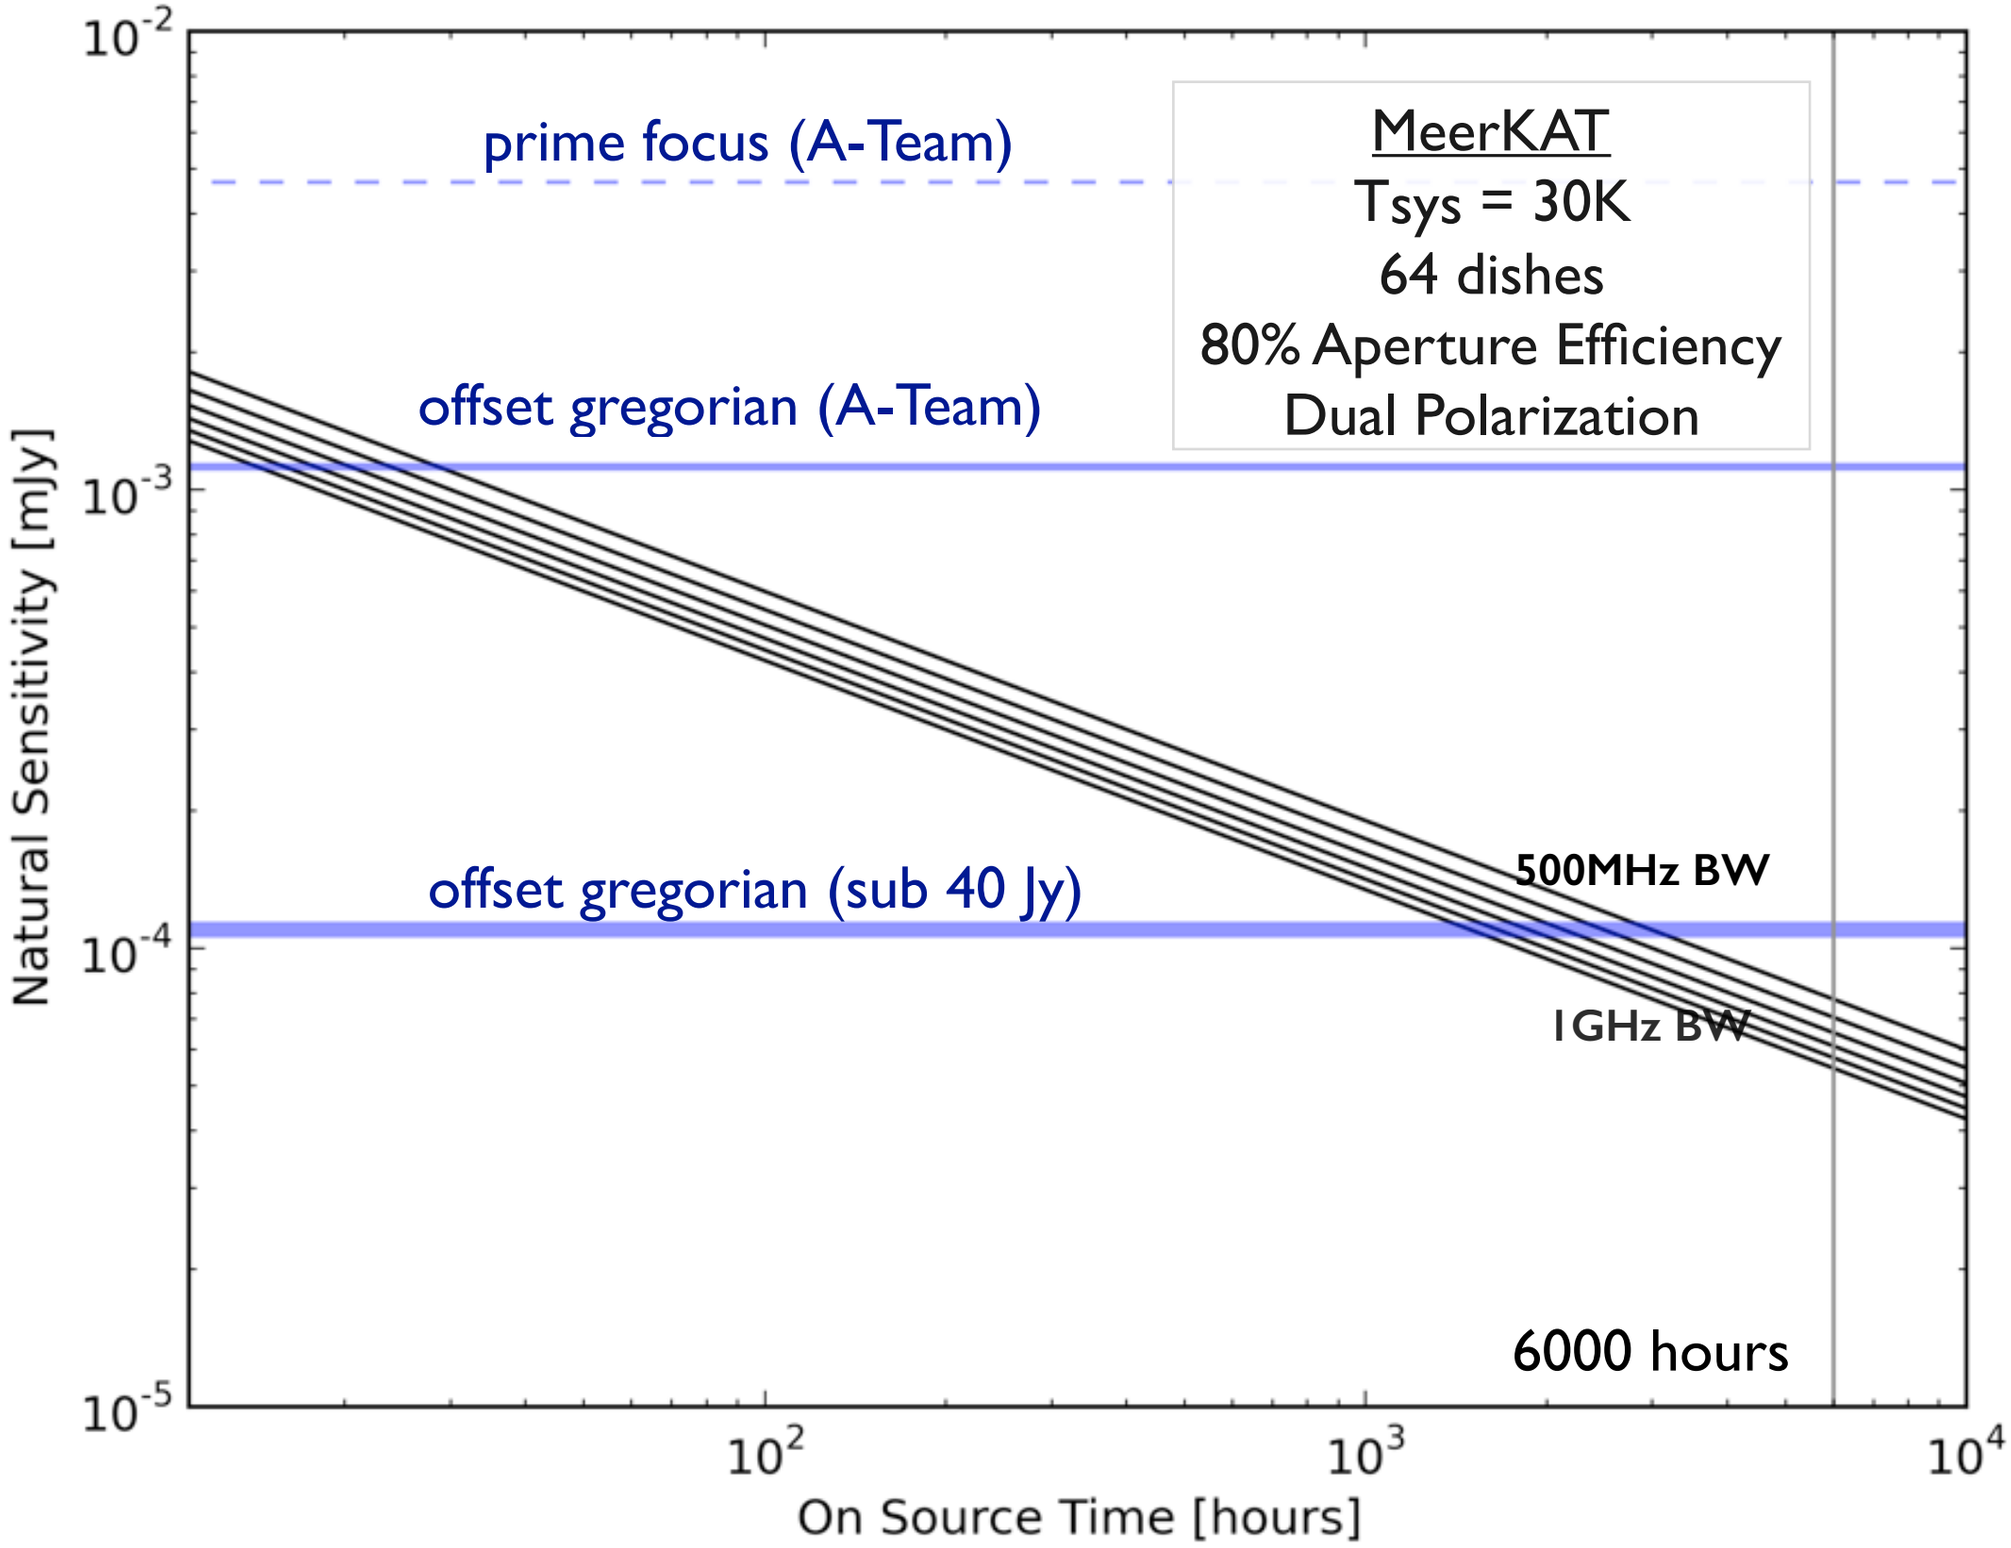
prime focus (A-Team)
offset gregorian (A-Team)
MeerKAT
Tsys = 30K
64 dishes
80% Aperture Efficiency
Dual Polarization
offset gregorian (sub 40 Jy)
500MHz BW
1GHz BW
6000 hours
Natural Sensitivity [mJy]
On Source Time [hours]
10^-2
10^-3
10^-4
10^-5
10^2
10^3
10^4

Sensitivity

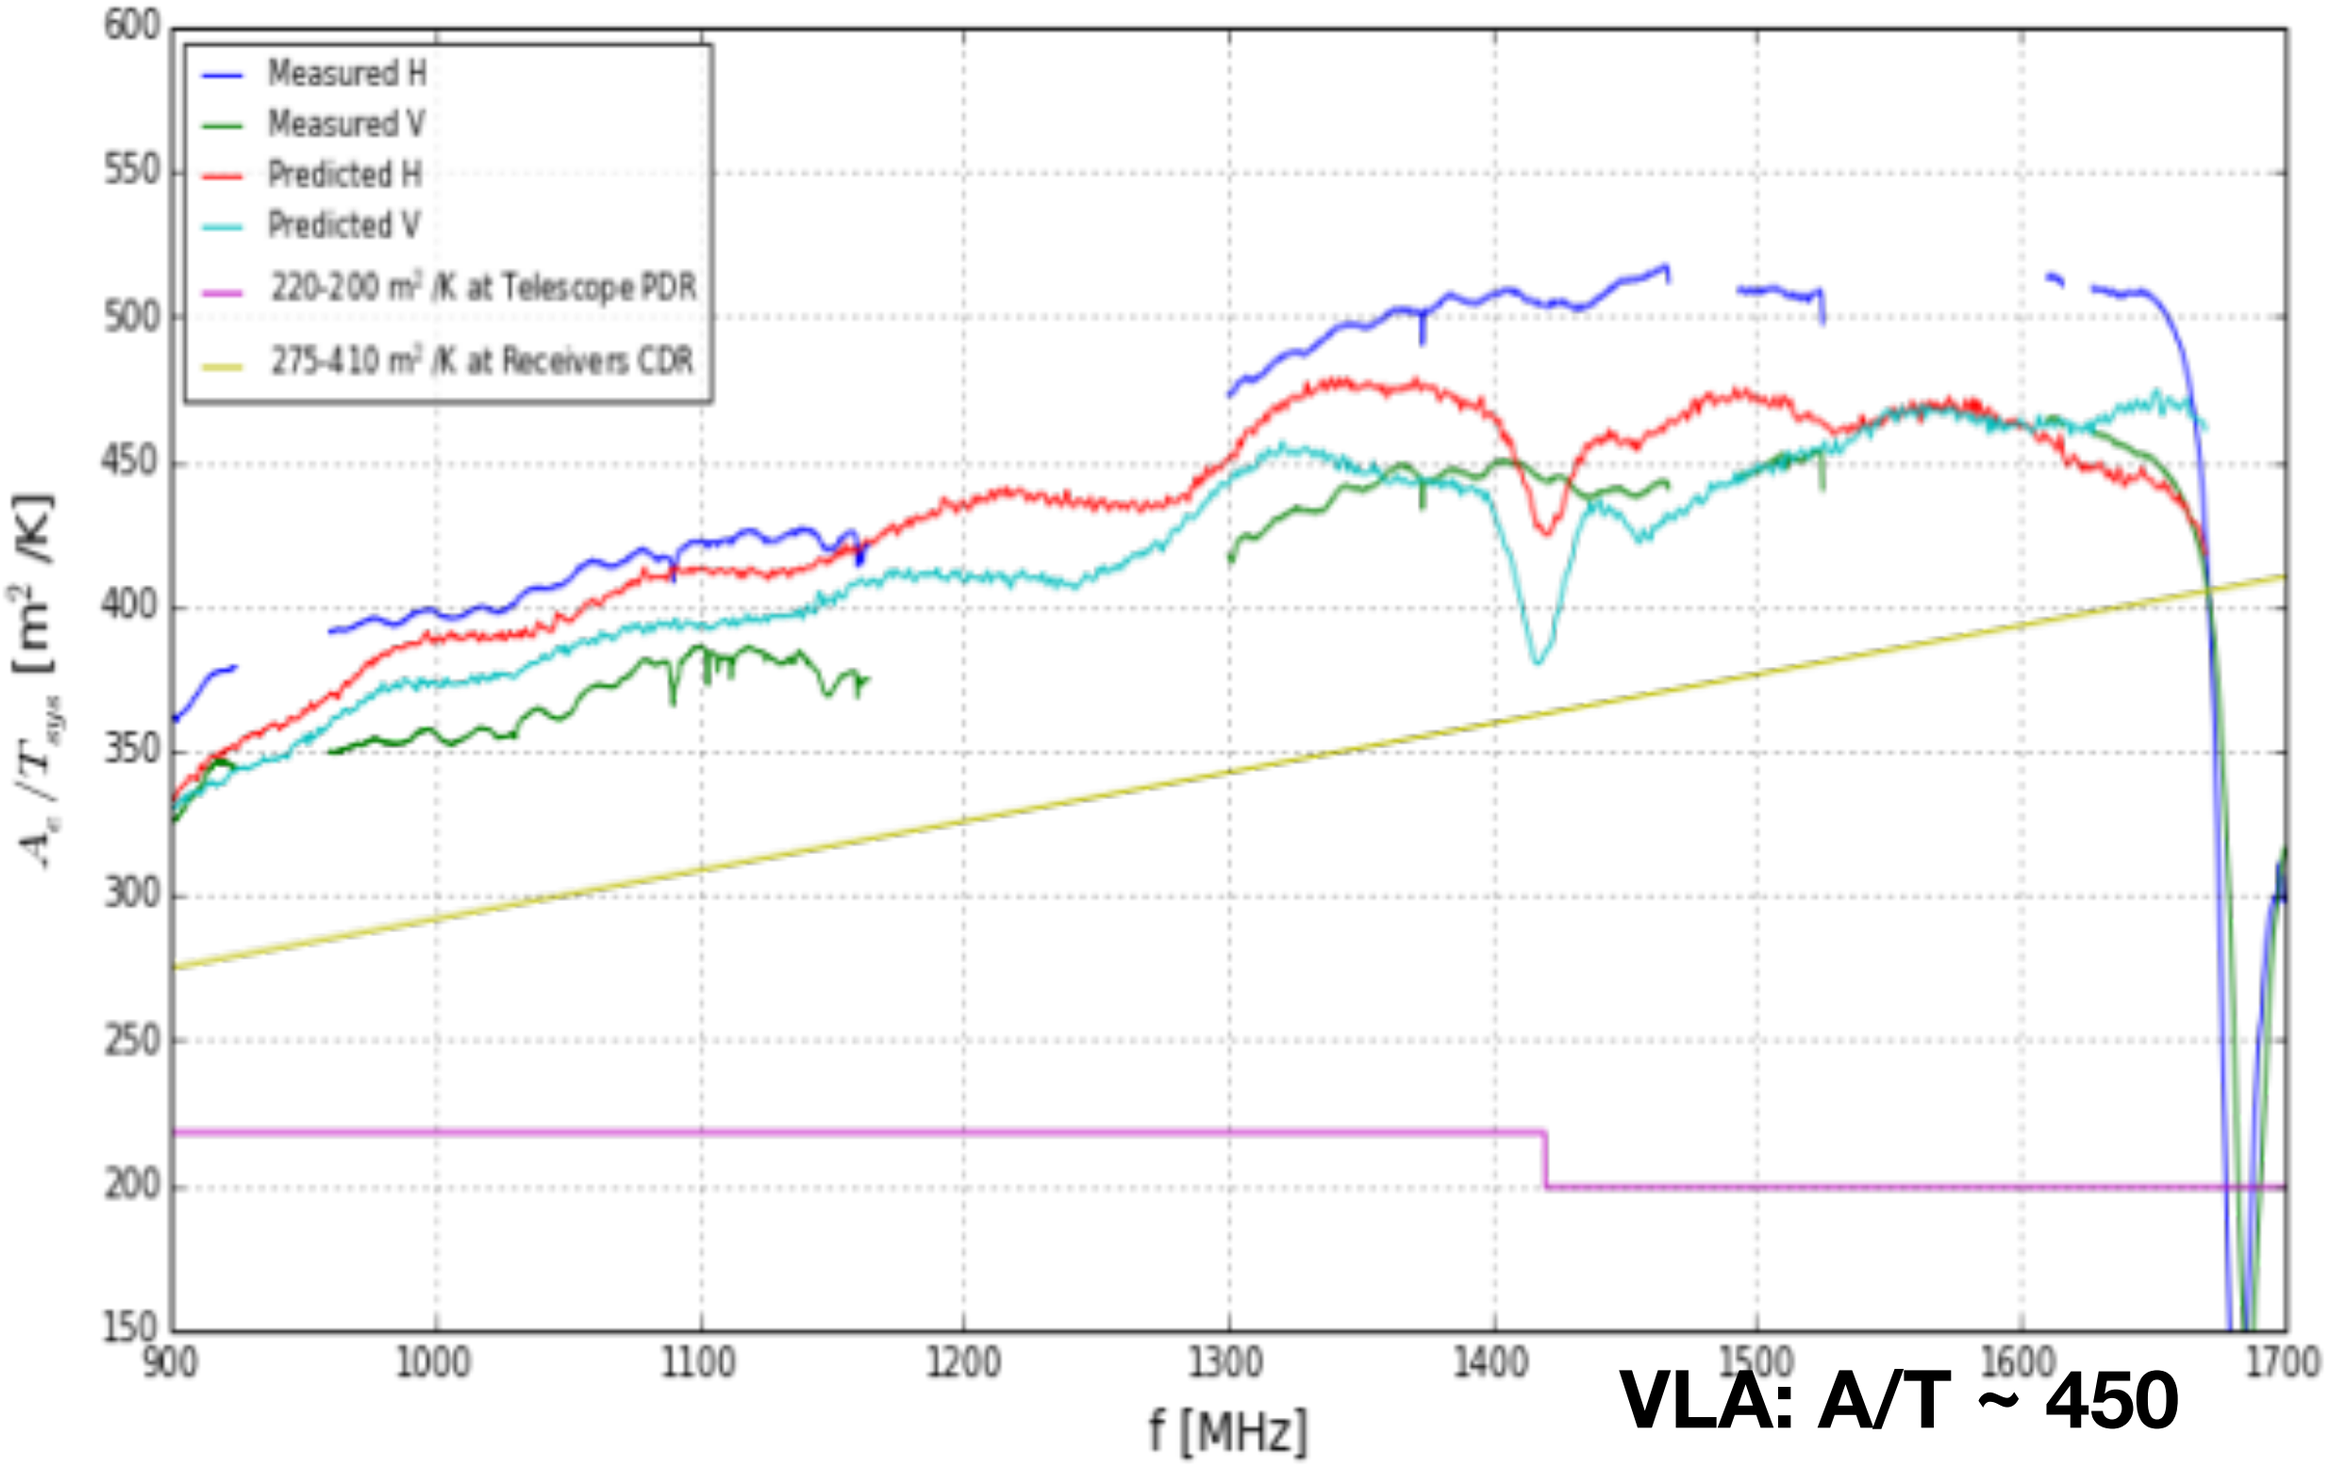

VLA: A/T ~ 450

Progress

- Currently AR1.5:
- 32-dishes with receivers installed.
    - ROACH2 Correlator: 32-inputs (16 dual polarization)
        - AR1.5 images and commissioning reports are available to the MLSPs.
- April 2018 AR3:
    - 64-dishes with receivers.
    - SKARAB Correlator.

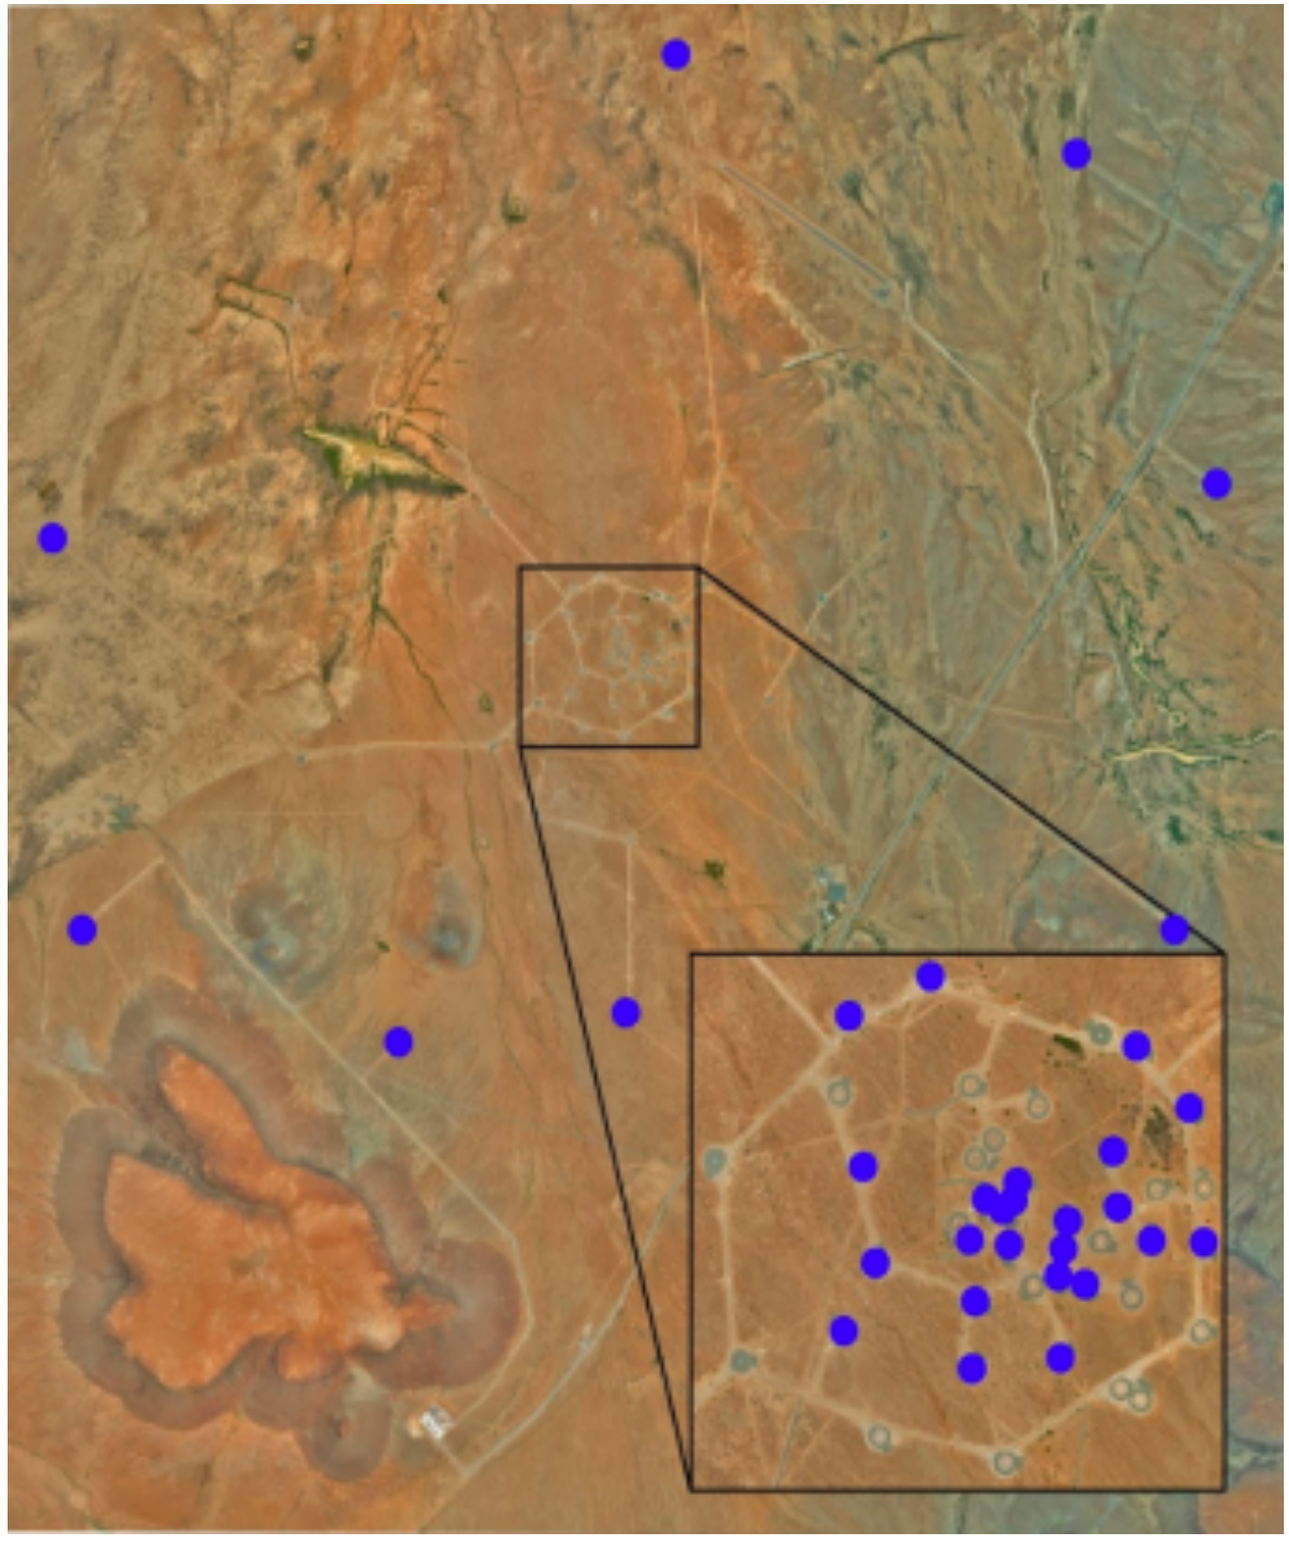

Timeline

- **AR-1: 16 receptor array**
  - First light: 30 June 2016
  - L-band receivers only
  - ROACH-2 correlator/ beamformer (CBF)
- **AR 2: 32 receptor array**
  - 32 single-polarisation inputs (ROACH-2 CBF): 31 March 2017
  - 32 dual-polarisation inputs (SKARAB CBF): late 2017
  - SKARAB CBF
- **AR 3: 64 receptor array**
  - Available by 31 March 2018
  - 64 dual-polarisation stations
  - SDP functionality
  - Hydrogen maser ensemble for time and frequency reference
- **ARs 4, 5 and 6**
  - Functionality to be implemented after AR3 in these releases includes:
  - Full sub-array support for commensal observations
  - Narrow band imaging mode

April 2017

IDIA Inter-University Institute for Data Intensive Astronomy

UNIVERSITY OF CAPE TOWN
IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD

ARC AFRICAN RESEARCH CLOUD

SKA SKA AFRICA

Correlator

	iBoB	ROACH	ROACH2	SNAP	SKARAB
Year Available	2005	2009	2010	2016	2016
Logic cells	53K	94K	476K	162-406K	693K
DSP slices	232	640	2016	600-1540	3600
BRAM capacity	4.2 Mb	8.8 Mb	38 Mb	11-28 Mb	53 Mb
SRAM capacity	2x18 Mb	2x36 Mb	4x144 Mb	-	-
SRAM bandwidth	9 Gb/s	43 Gb/s	200 Gb/s	-	-
DDR capacity (max)	-	1x8 Gb	1x16 Gb	-	-
DDR bandwidth	-	38 Gb/s	50 Gb/s	-	-
HMC capacity	-	-		-	<8x32 Gb
HMC bandwidth	-	-		-	<8x30 Gb/s
Ethernet ports	2x10 GbE	4x10 GbE	8x10 GbE	2x10 GbE	<16x40 GbE
ADC/DAC support	2xZDOK	2xZDOK	2xZDOK	1xZDOK, 3xHMCAD1511	4xMegarray

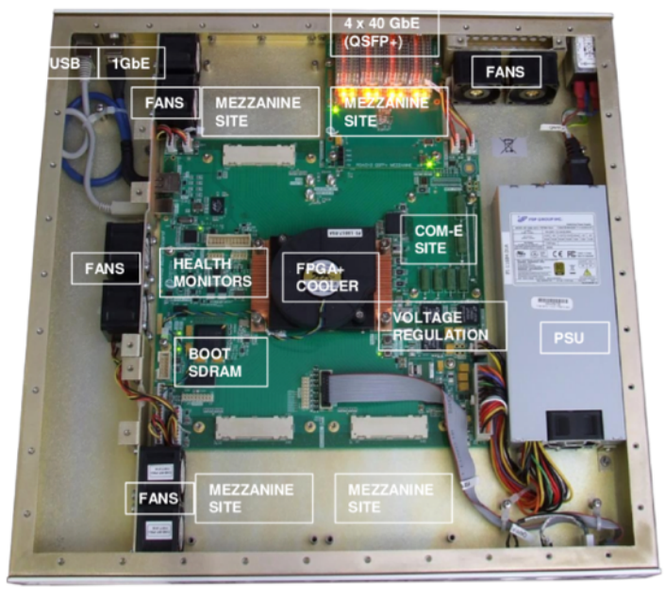

(a) The SKARAB platform is a modular processing platform based around a Xilinx Virtex 7 FPGA with expandible mezzanine slots for memory and Ethernet ports.

DDR to HMC

Hybrid Memory Cube uses stacked silicon and high speed serial interfaces — add independent memory interfaces to each processing node without running out of pins on the FPGA's.

Hickish et al 2016, A Decade of Developing Radio-Astronomy Instrumentation using CASPER Open-Source Technology

MeerKAT AR1.5 Images

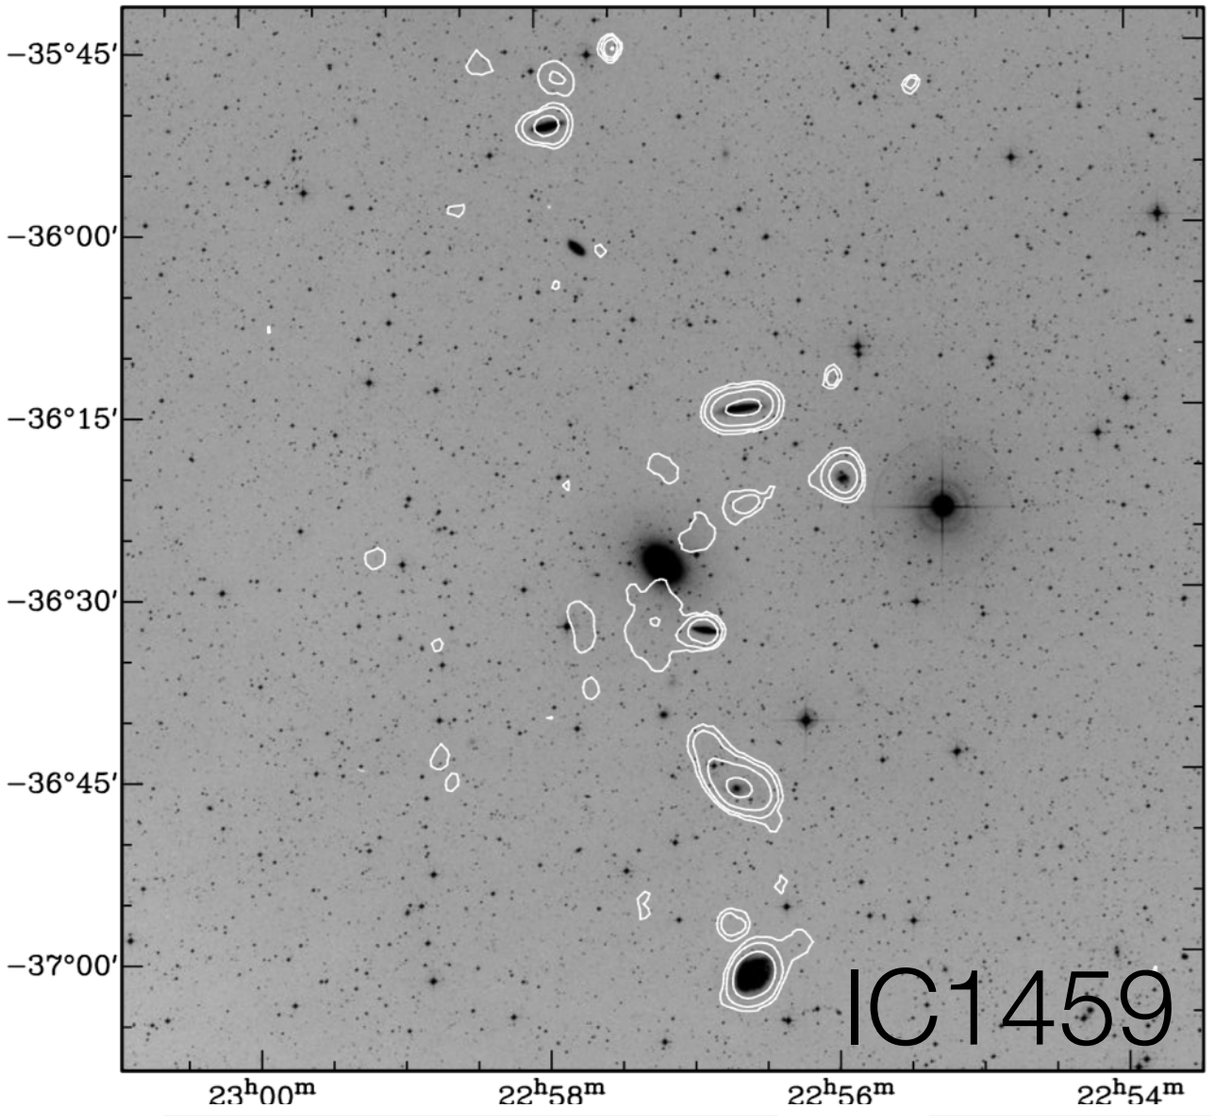

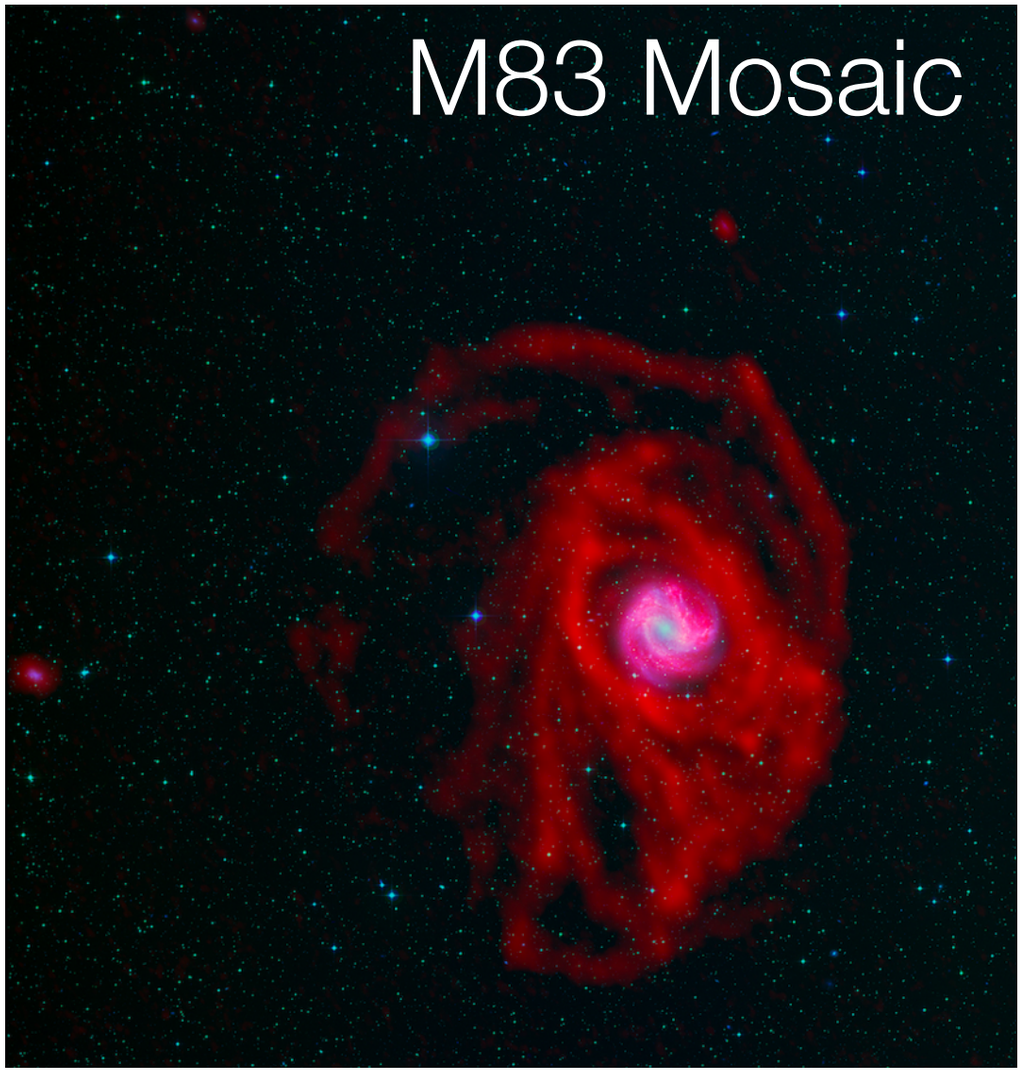

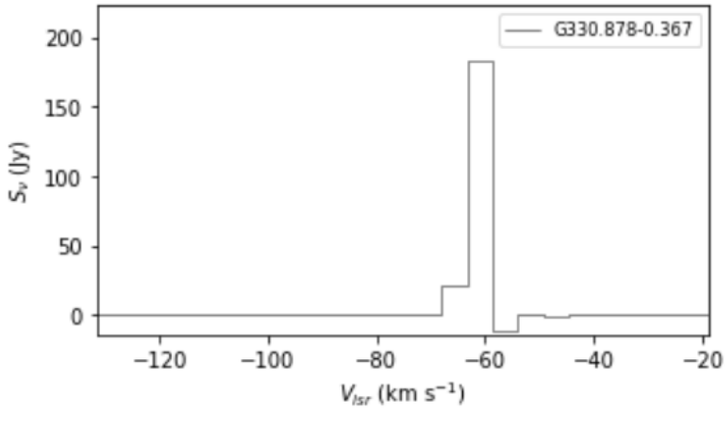

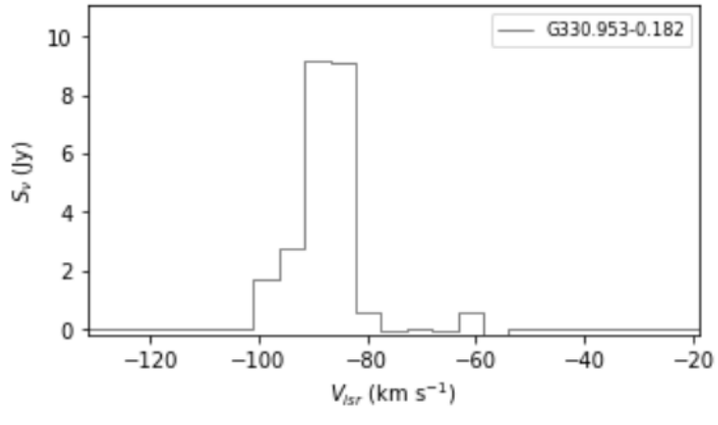

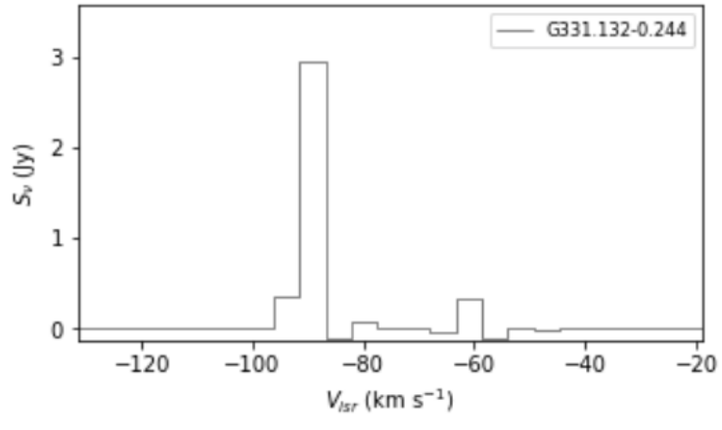

IDIA for Data Intensive Astronomy

UNIVERSITY OF CAPE TOWN
IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD

AARC AFRICAN RESEARCH CLOUD

SKA AFRICA

The MeerKAT Team & Paolo Serra

HIPASS J0358-81 with MeerKAT

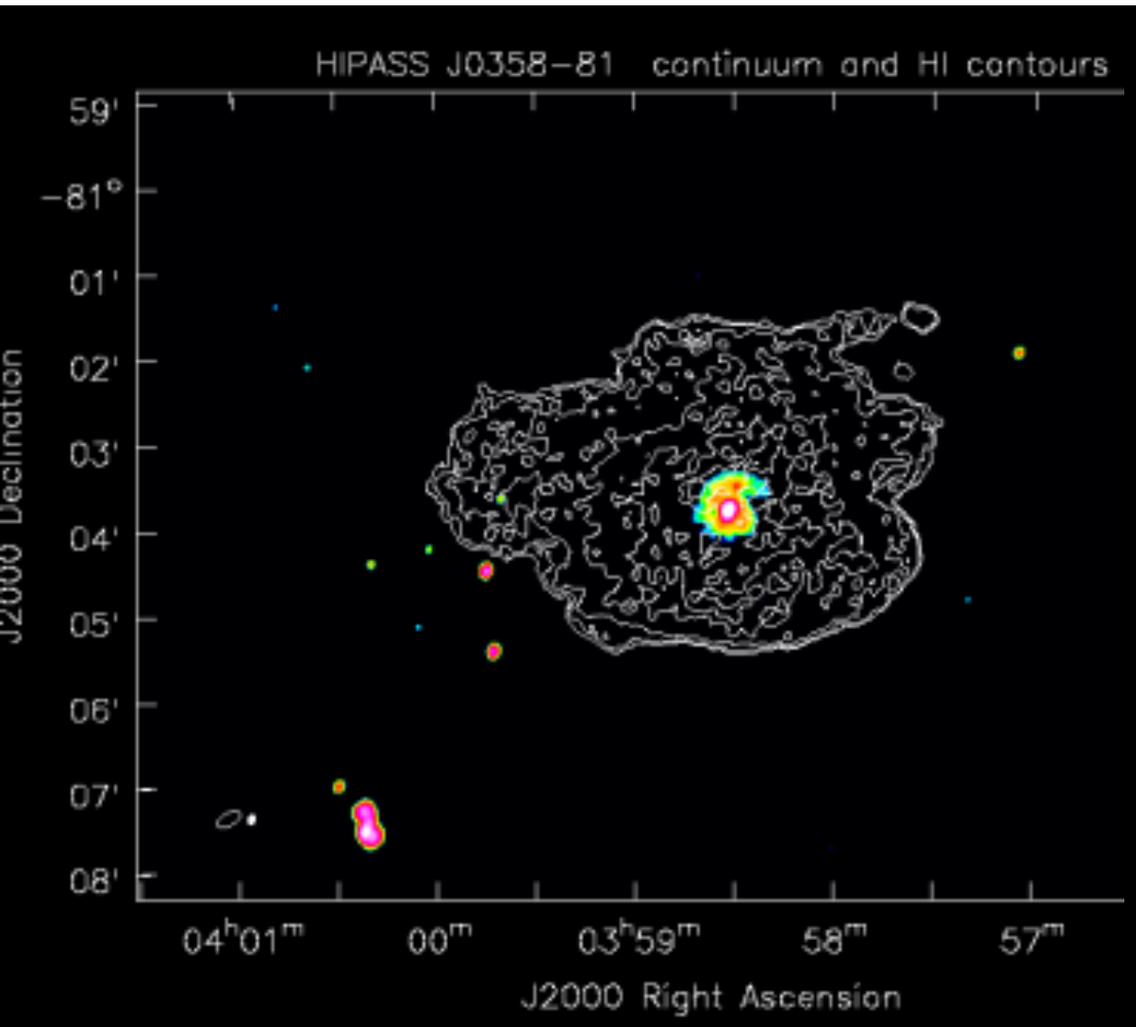

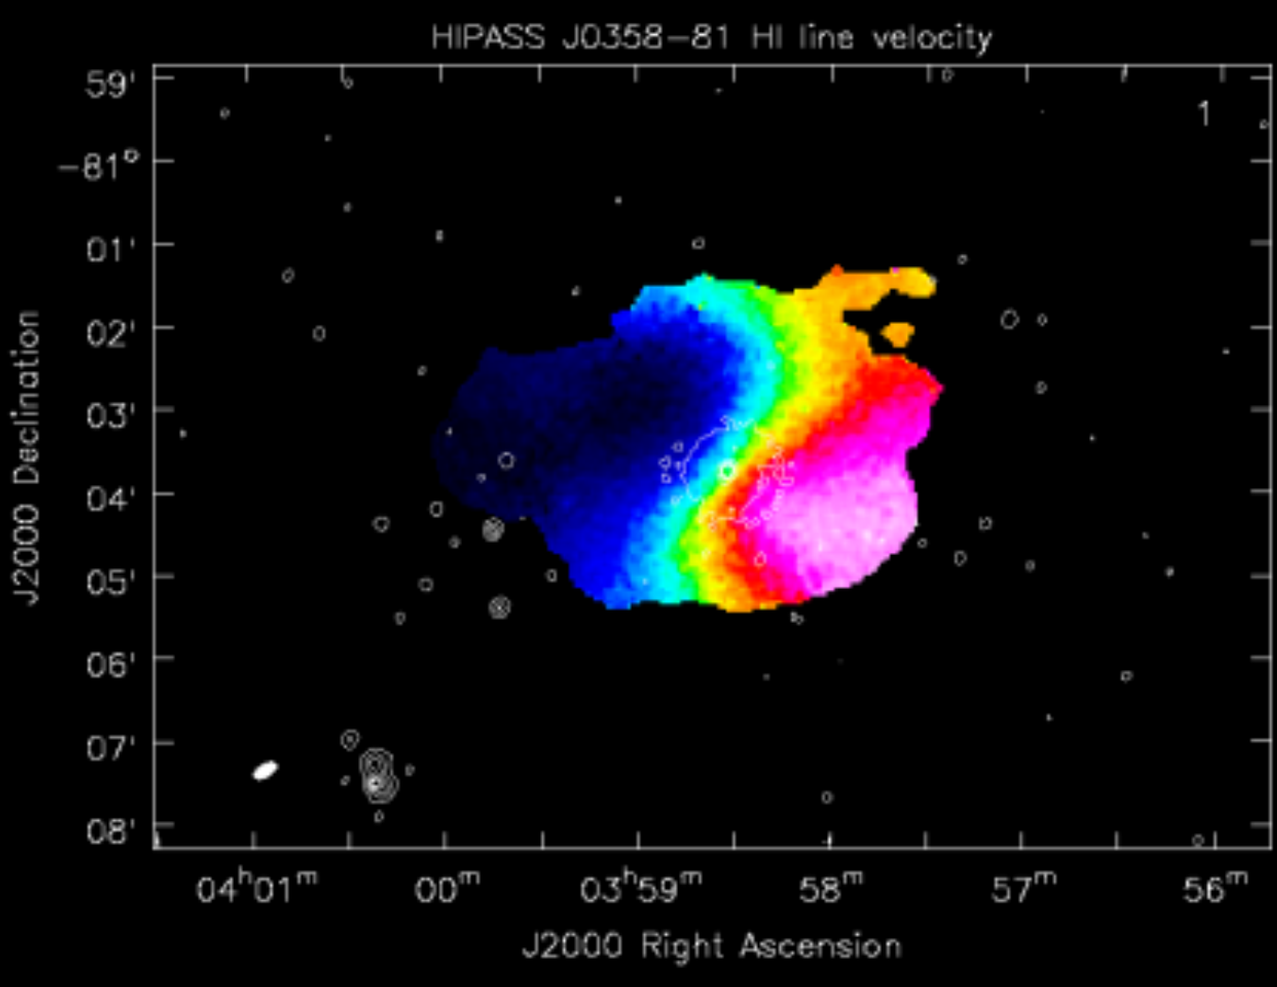

The MeerKAT Team

IDIA Inter-University Institute for Data Intensive Astronomy

UNIVERSITY OF CAPE TOWN
IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD

ARC AFRICAN RESEARCH CLOUD

SKA AFRICA

MeerKAT Large Science Projects

Priority Group 1

Radio Pulsar Timing: Testing Einstein's theory of gravity and gravitational radiation - Investigating the physics of enigmatic neutron stars through observations of pulsars. **7860**

LADUMA (Looking at the Distant Universe with the MeerKAT Array) - An ultra-deep survey of neutral hydrogen gas in the early universe. **5000**

Priority Group 2

MESMER (MeerKAT Search for Molecules in the Epoch of Re-ionisation) - Searching for CO at high red-shift (z>7) to investigate the role of molecular hydrogen in the early universe. **6500**

MeerKAT Absorption Line Survey for atomic hydrogen and OH lines in absorption against distant continuum sources. **4000**

MHONGOOSE (MeerKAT HI Observations of Nearby Galactic Objects: Observing Southern Emitters) - Investigations of different types of galaxies; dark matter and the cosmic web. **6000**

TRAPUM (Transients and Pulsars with MeerKAT) - Searching for, and investigating new and exotic pulsars. **3000**

A MeerKAT HI Survey of the Fornax Cluster (Galaxy formation and evolution in the cluster environment). **2450**

MeerGAL (MeerKAT High Frequency Galactic Plane Survey) - Galactic structure and dynamics, distribution of ionised gas, recombination lines, interstellar molecular gas and masers. **3300**

MIGHTEE (MeerKAT International GigaHertz Tiered Extragalactic Exploration Survey) - Deep continuum observations of the earliest radio galaxies. **1950**

ThunderKAT (The Hunt for Dynamic and Explosive Radio Transients with MeerKAT) - eg gamma ray bursts, novae and supernovae, plus new types of transient radio sources. **3000**

Review of MeerKAT Large Projects

- A: high impact science that is uniquely well-matched to the capabilities of MeerKAT
- B-ranked: excellent science, but technically challenging.

Large Survey Project (LSP) Components	*Requested Time (hrs)*	*Readers' Science Score*	*Panel Ranking*
MeerTime (binary)	1440	3.87	A
MHONGOOSE	1650	3.55	A
MeerTime (MSPs)	2160	3.58	A
LADUMA	3424	3.84	A
Fornax	900	3.41	A
TRAPUM (Fermi sources)	338	3.60	A
MeerTime (1000 PTA)	720	3.78	A/B
ThunderKAT (CVs)	250	3.42	B
MIGHTEE (L band)	979	3.23	B
ThunderKAT (GRBs)	330	3.42	B
MeerTime (GCs)	1080	3.38	B
MALS (UHF band)	2320	N/A	B
TRAPUM (nearby galaxies)	226	3.28	B
TRAPUM (GCs)	320	3.22	B
TRAPUM (SNR, PWN, TeV)	92	N/A	B
ThunderKAT (SNe Ia)	200	3.08	B
MIGHTEE (S band)	948	2.77	B/C

MIGHTEE
Taylor & Jarvis

- The MeerKAT International GHz Tuned Extragalactic Exploration.
  - Originally tiered, without dedicated HI component.
  - Longer baselines to help with confusion.
- 2016/17 Revision
  1. Number of pointings for effective survey area.
  2. Updated instrumental thermal noise.
  3. Increase in depth due to overlapping pointings.
  4. Confusion noise.

MIGHTEE Survey Design

- Sensitivity
  - Originally: 1μJy/beam, 1500-h.
  - Revised: 1606-h.
    - L-band: 2μJy/beam, 816-h (confusion — Wilman et al. 2008, Condon et al. 2012).
    - S-band: 1μJy/beam, 790-h.

- Footprint
  - Original: 35 sq. deg.
  - Revised: ~20 sq. deg. effective in L-band, ~6 sq. deg. in S-band.
    - XMM-LSS, 6.7 sq. deg (L-band only).
    - E-CDFS, 8.3 sq. deg (L-band, S-band: 4 sq. deg.), commensal with LADUMA.
    - ELAIS-S1, 1.6 sq. deg (L-band only).
    - COSMOS, ~1 sq. deg. (L-band, S-band: 1.5 sq. deg.)

MIGHTEE Footprint

XMM-LSS

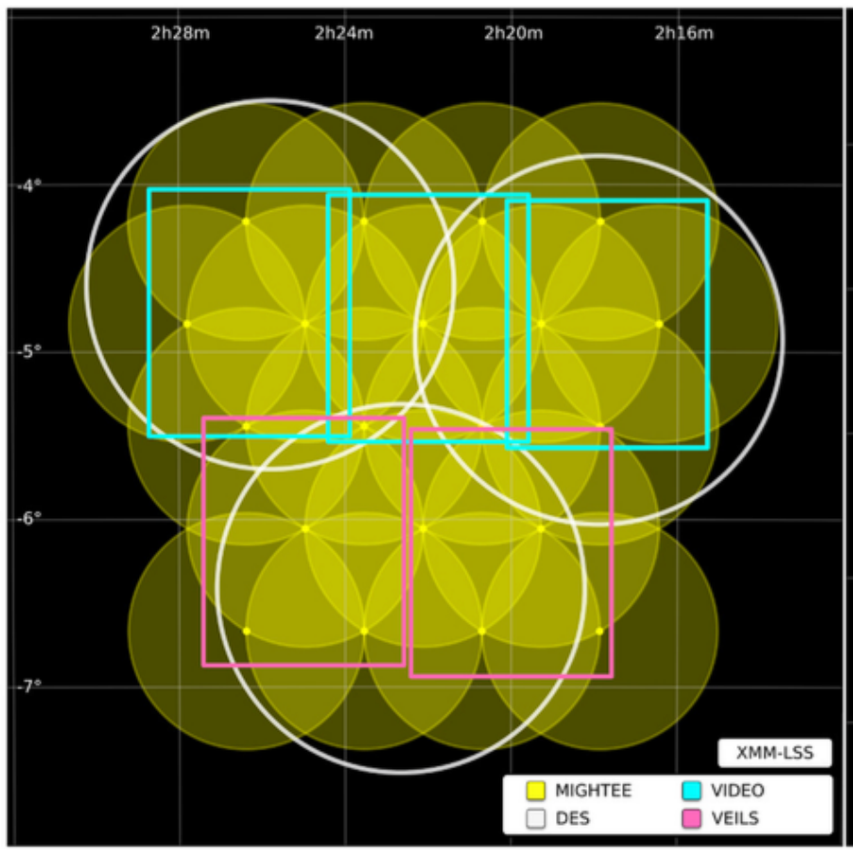

E-CDFS

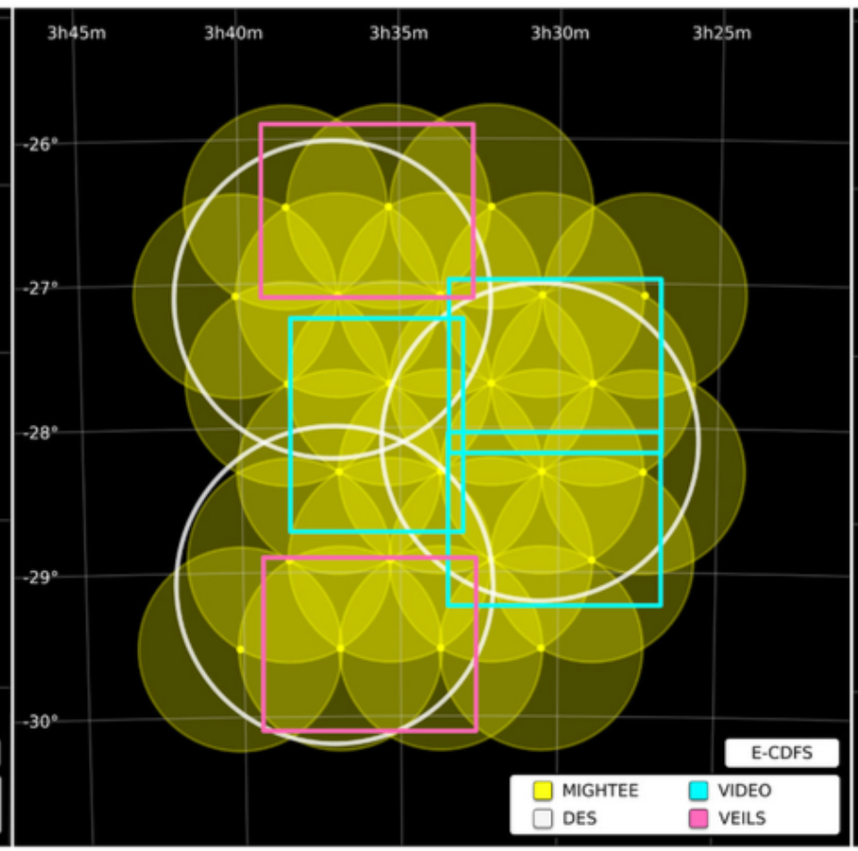

ELAIS-S1

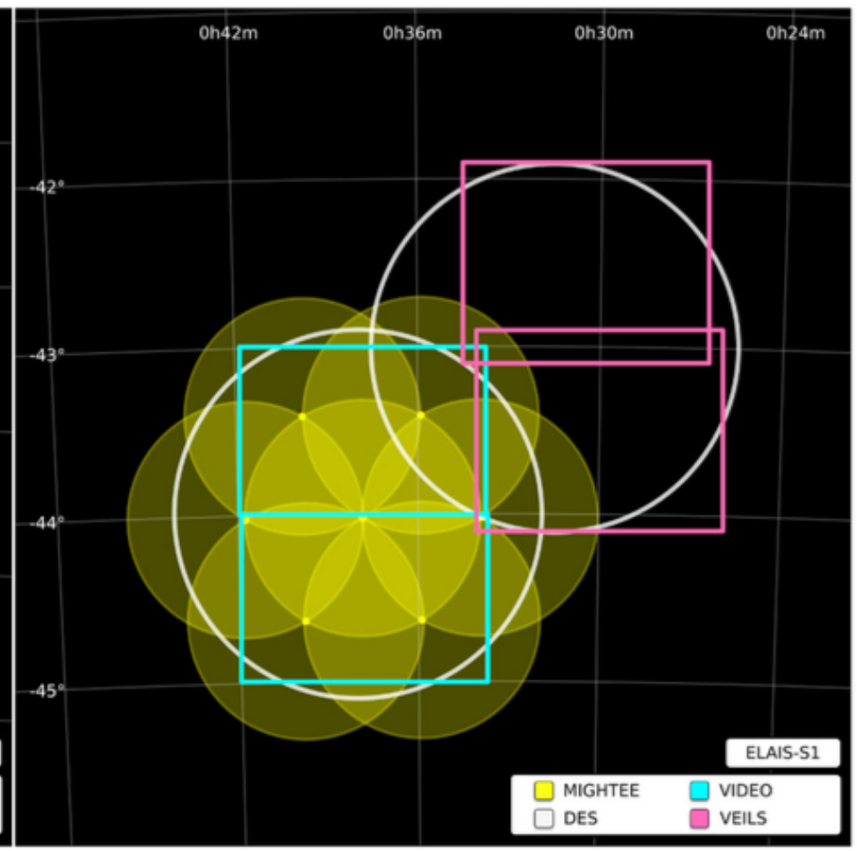

+COSMOS

IDIA Inter-University Institute for Data Intensive Astronomy

UNIVERSITY OF CAPE TOWN
IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD

ARC AFRICAN RESEARCH CLOUD

SKA AFRICA

MIGHTEE Science

- Quenching — the role of environment.
- AGN polarisation as a probe of environment.
- AGN fuelling and feedback.
- Tracing the mechanical feedback from AGN jets.
- HI as a direct probe of neutral gas accretion and feedback in AGN.
- Large scale structure and cosmology.
- Evolution of bias.
- Weak lensing.
- The Magnetic Cosmic Web.
- Resolving Massive Galaxies out to Cosmological Distances.
- Low-mass nearby galaxies (missing satellites).
- Clusters and Magnetic Fields.
- Emergence/evolution of Magnetic Fields in Galaxies.
- Stacking.

MIGHTEE Science

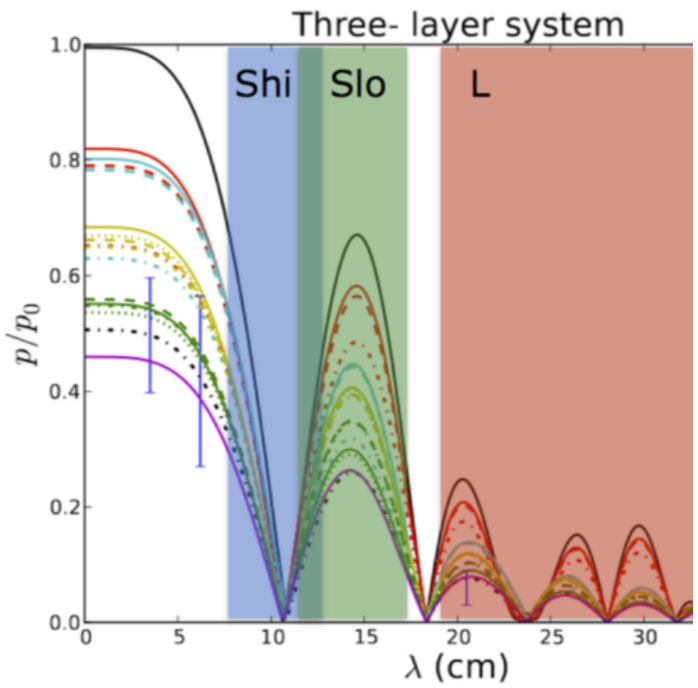

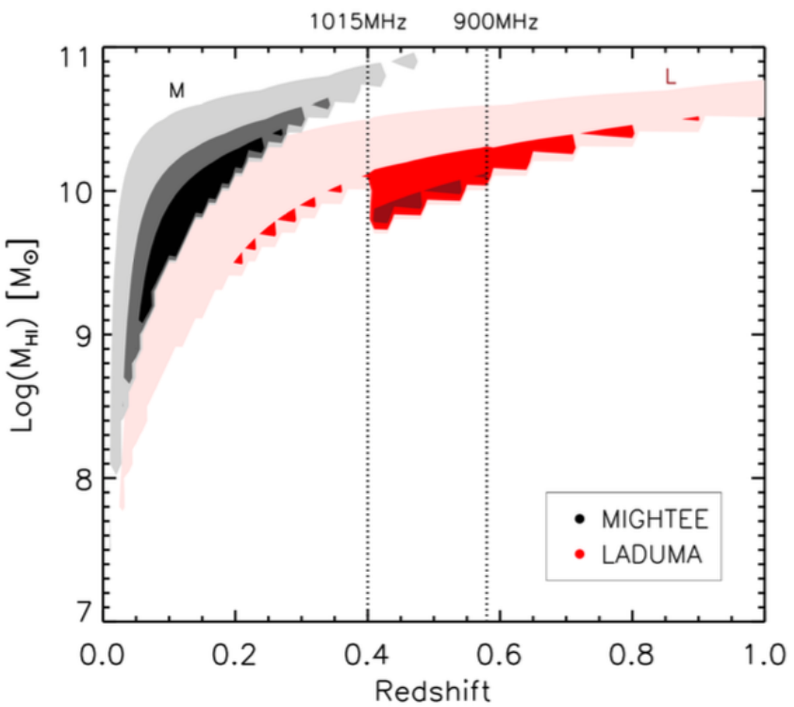

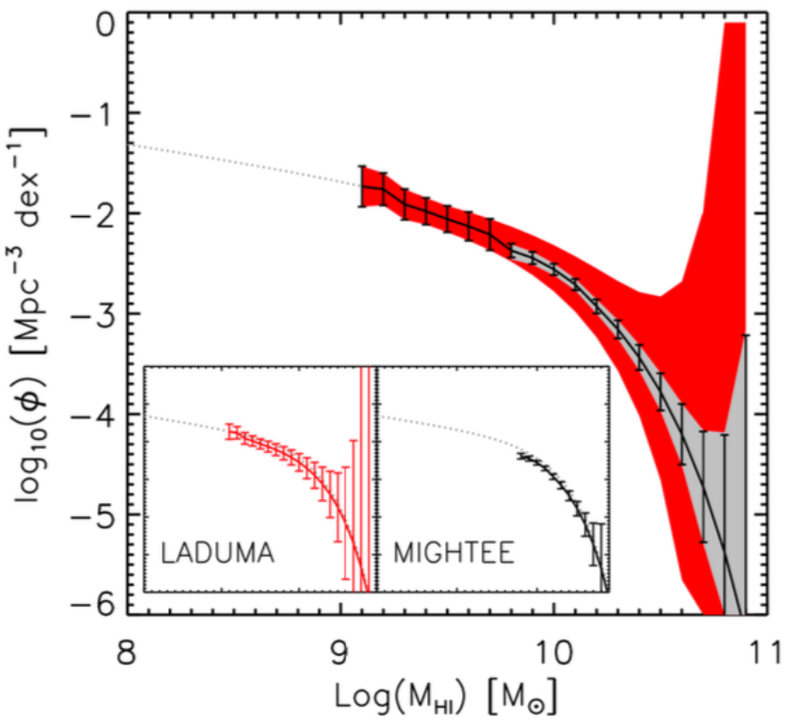

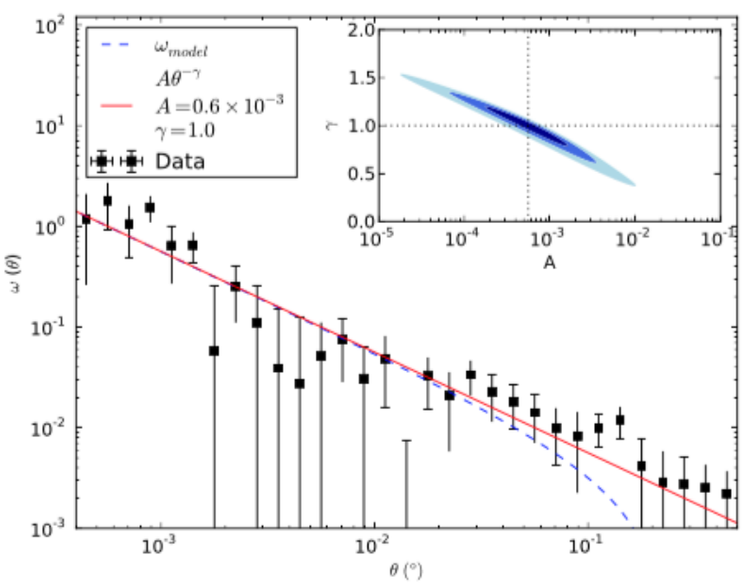

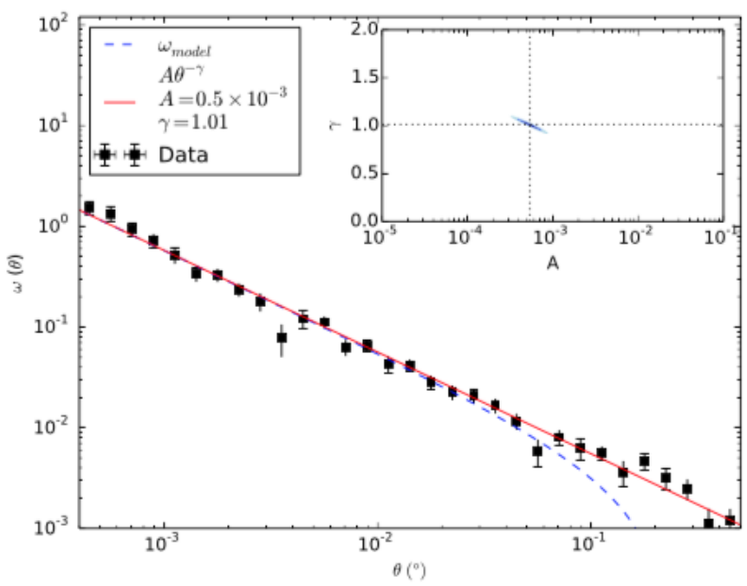

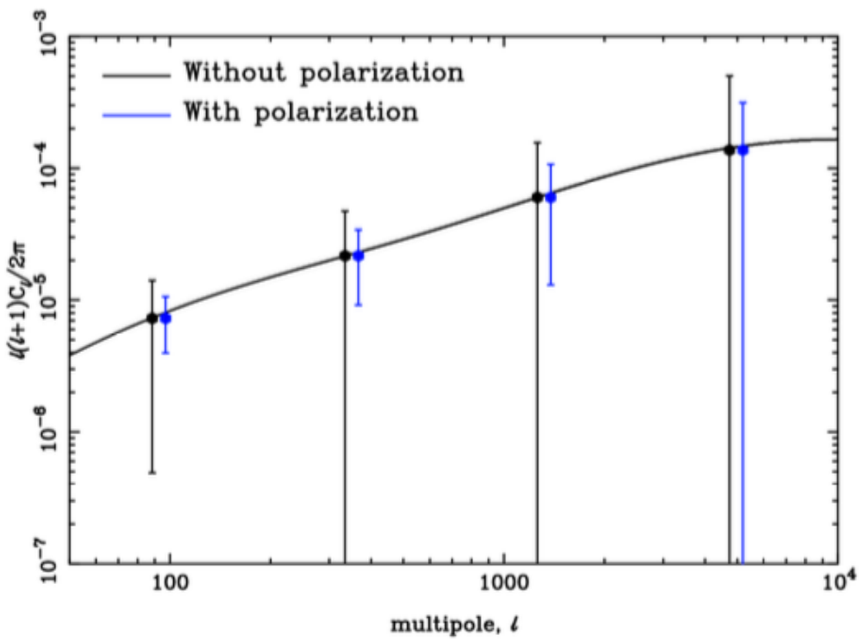

uGMRT-MIGHTEE

- Proposal: Pilot Project with Upgraded GMRT (see Dharam Vir Lal's Talk later), Taylor et al.
  - Resolution: $\Theta_{GMRT@600MHz} \sim \Theta_{MeerKAT@1400MHz}$
  - E-CDFS: 3x30-h pointings.
  - 550–850 MHz uGMRT band (band 4)
  - Theoretical: 2μJy/beam (MIGHTEE).
  - MIGHTEE/LADUMA partnership, exploiting complementarity.

uGMRT-MIGHTEE

- Science
  - Star Formation History.
  - Faint AGN.
  - Steep Spectrum Objects.
  - Cosmic Magnetic Fields.
  - Atomic Hydrogen(!).

Ocran et al. 2017, Ocran et al. (in prep)

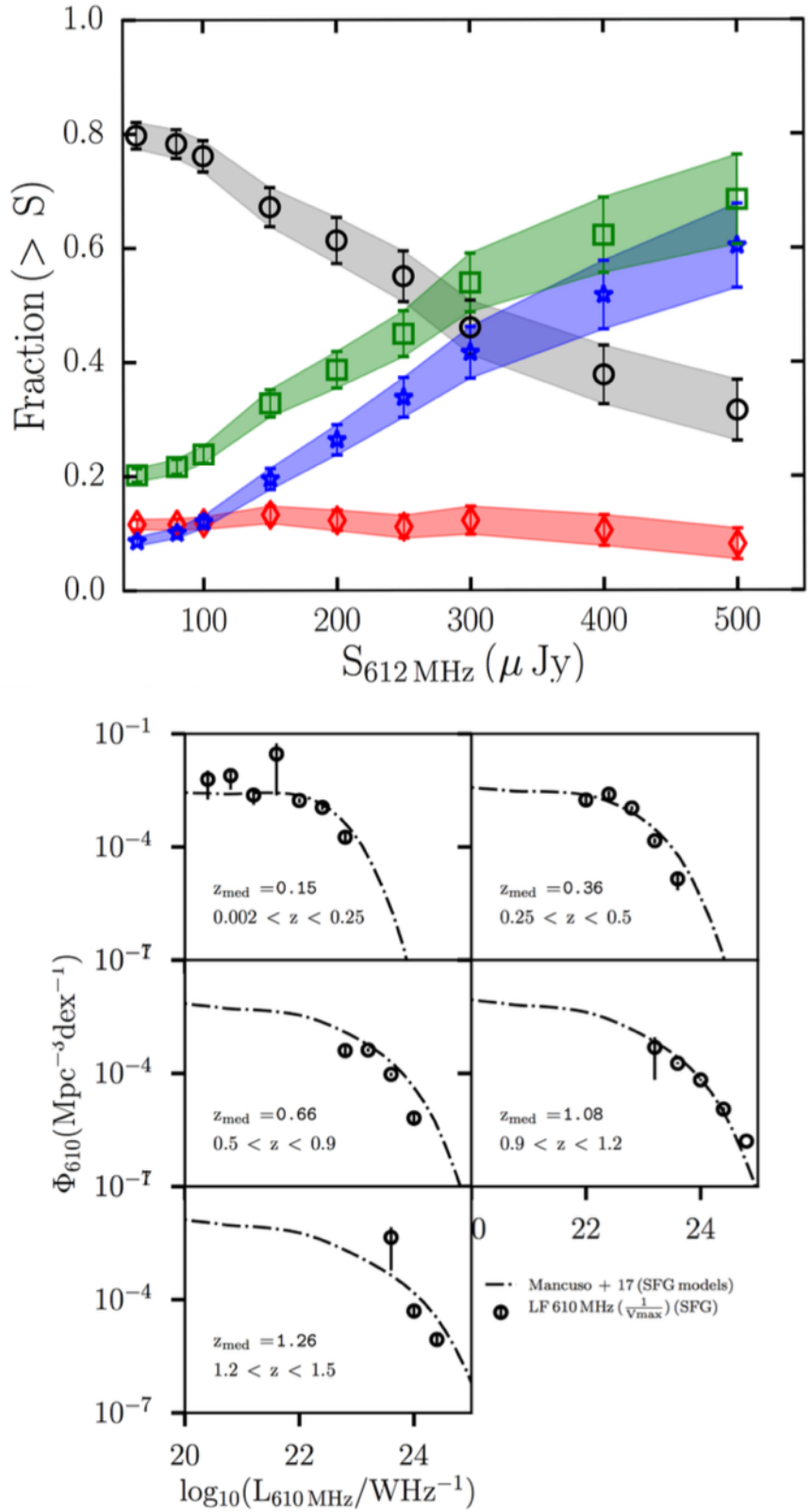

uGMRT-MIGHTEE

- Technical issues in prep for full uGMRT-MIGHTEE proposal:
  - DR, Confusion, optimal imaging for resolution/noise.
  - Commissioning and development of cloud-based pipelines (IDIA): calibration, imaging and analysis.
  - Joint deconvolution of uGMRT and MeerKAT data: ultra-wideband MFS and RM Sythesis(!).
  - Limits of calibratability and instrumental effects.
  - High spatial/spectral resolution modelling of PKS0326-268 for off-axis correction.

MeerKAT (HI) Fornax Survey
PI: Paolo Serra

- Fornax Cluster: closest cluster visible in the south, 100 x 9-h, 12 sq. deg.
  1. Detect the tails of HI gas removed from galaxies living in Fornax.
  2. Study the HI mass function in Fornax.
  3. Physical processes driving galaxy evolution in clusters.
  4. Detect the faint HI gas which should be found between galaxies in the cosmic web.

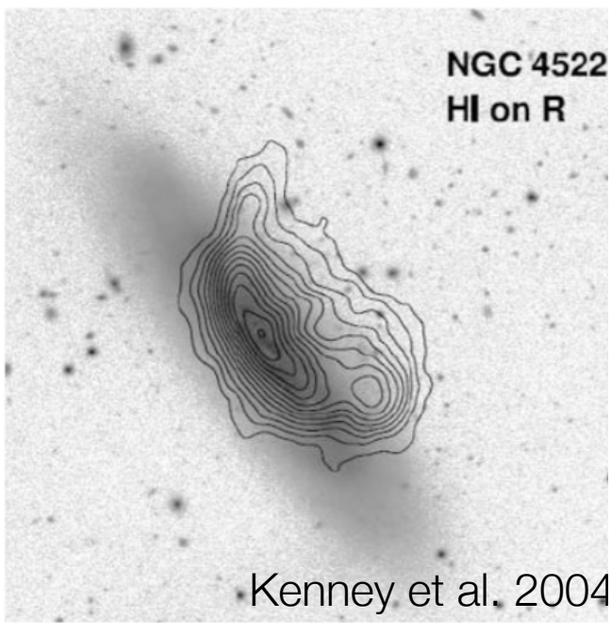

Kenney et al. 2004

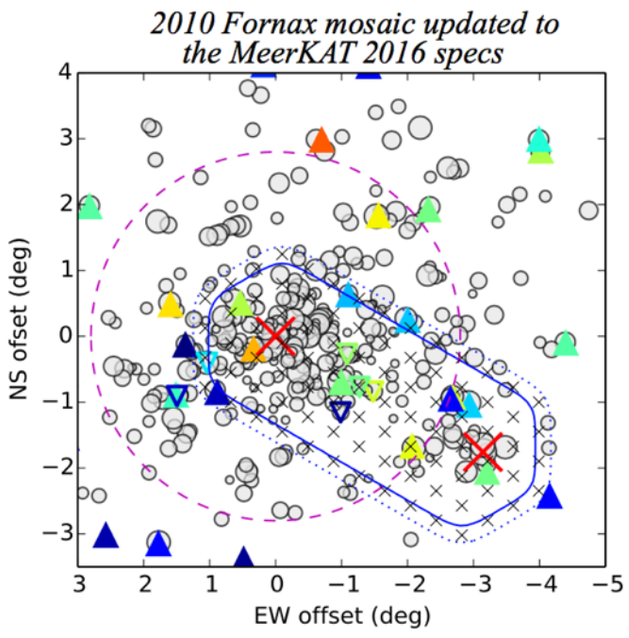

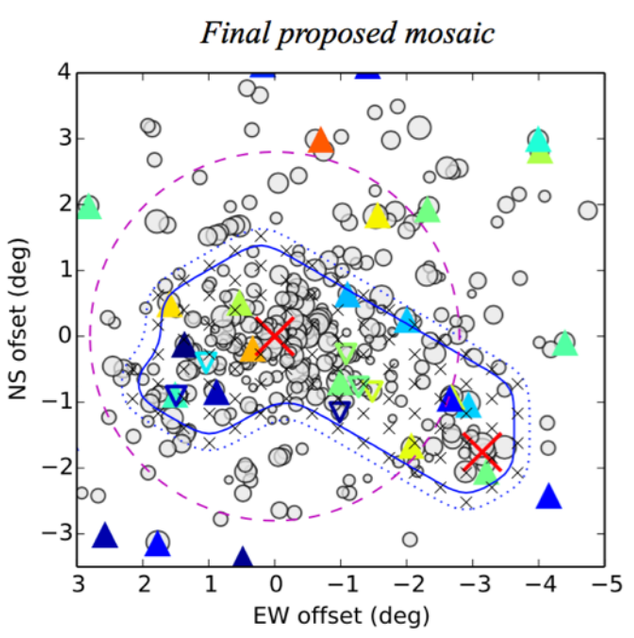

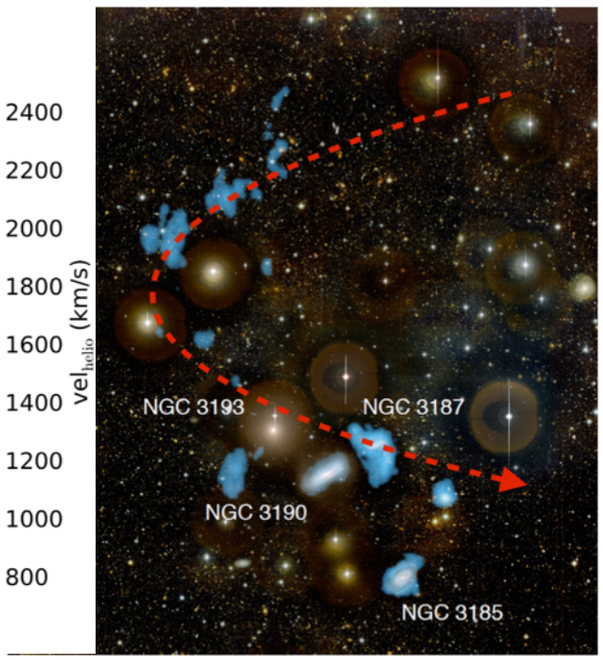

HCG 44, Serra et al. 2013.

IDIA Inter-University Institute for Data Intensive Astronomy
UNIVERSITY OF CAPE TOWN IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD
ARC
SKA AFRICA

External (to site) data flows

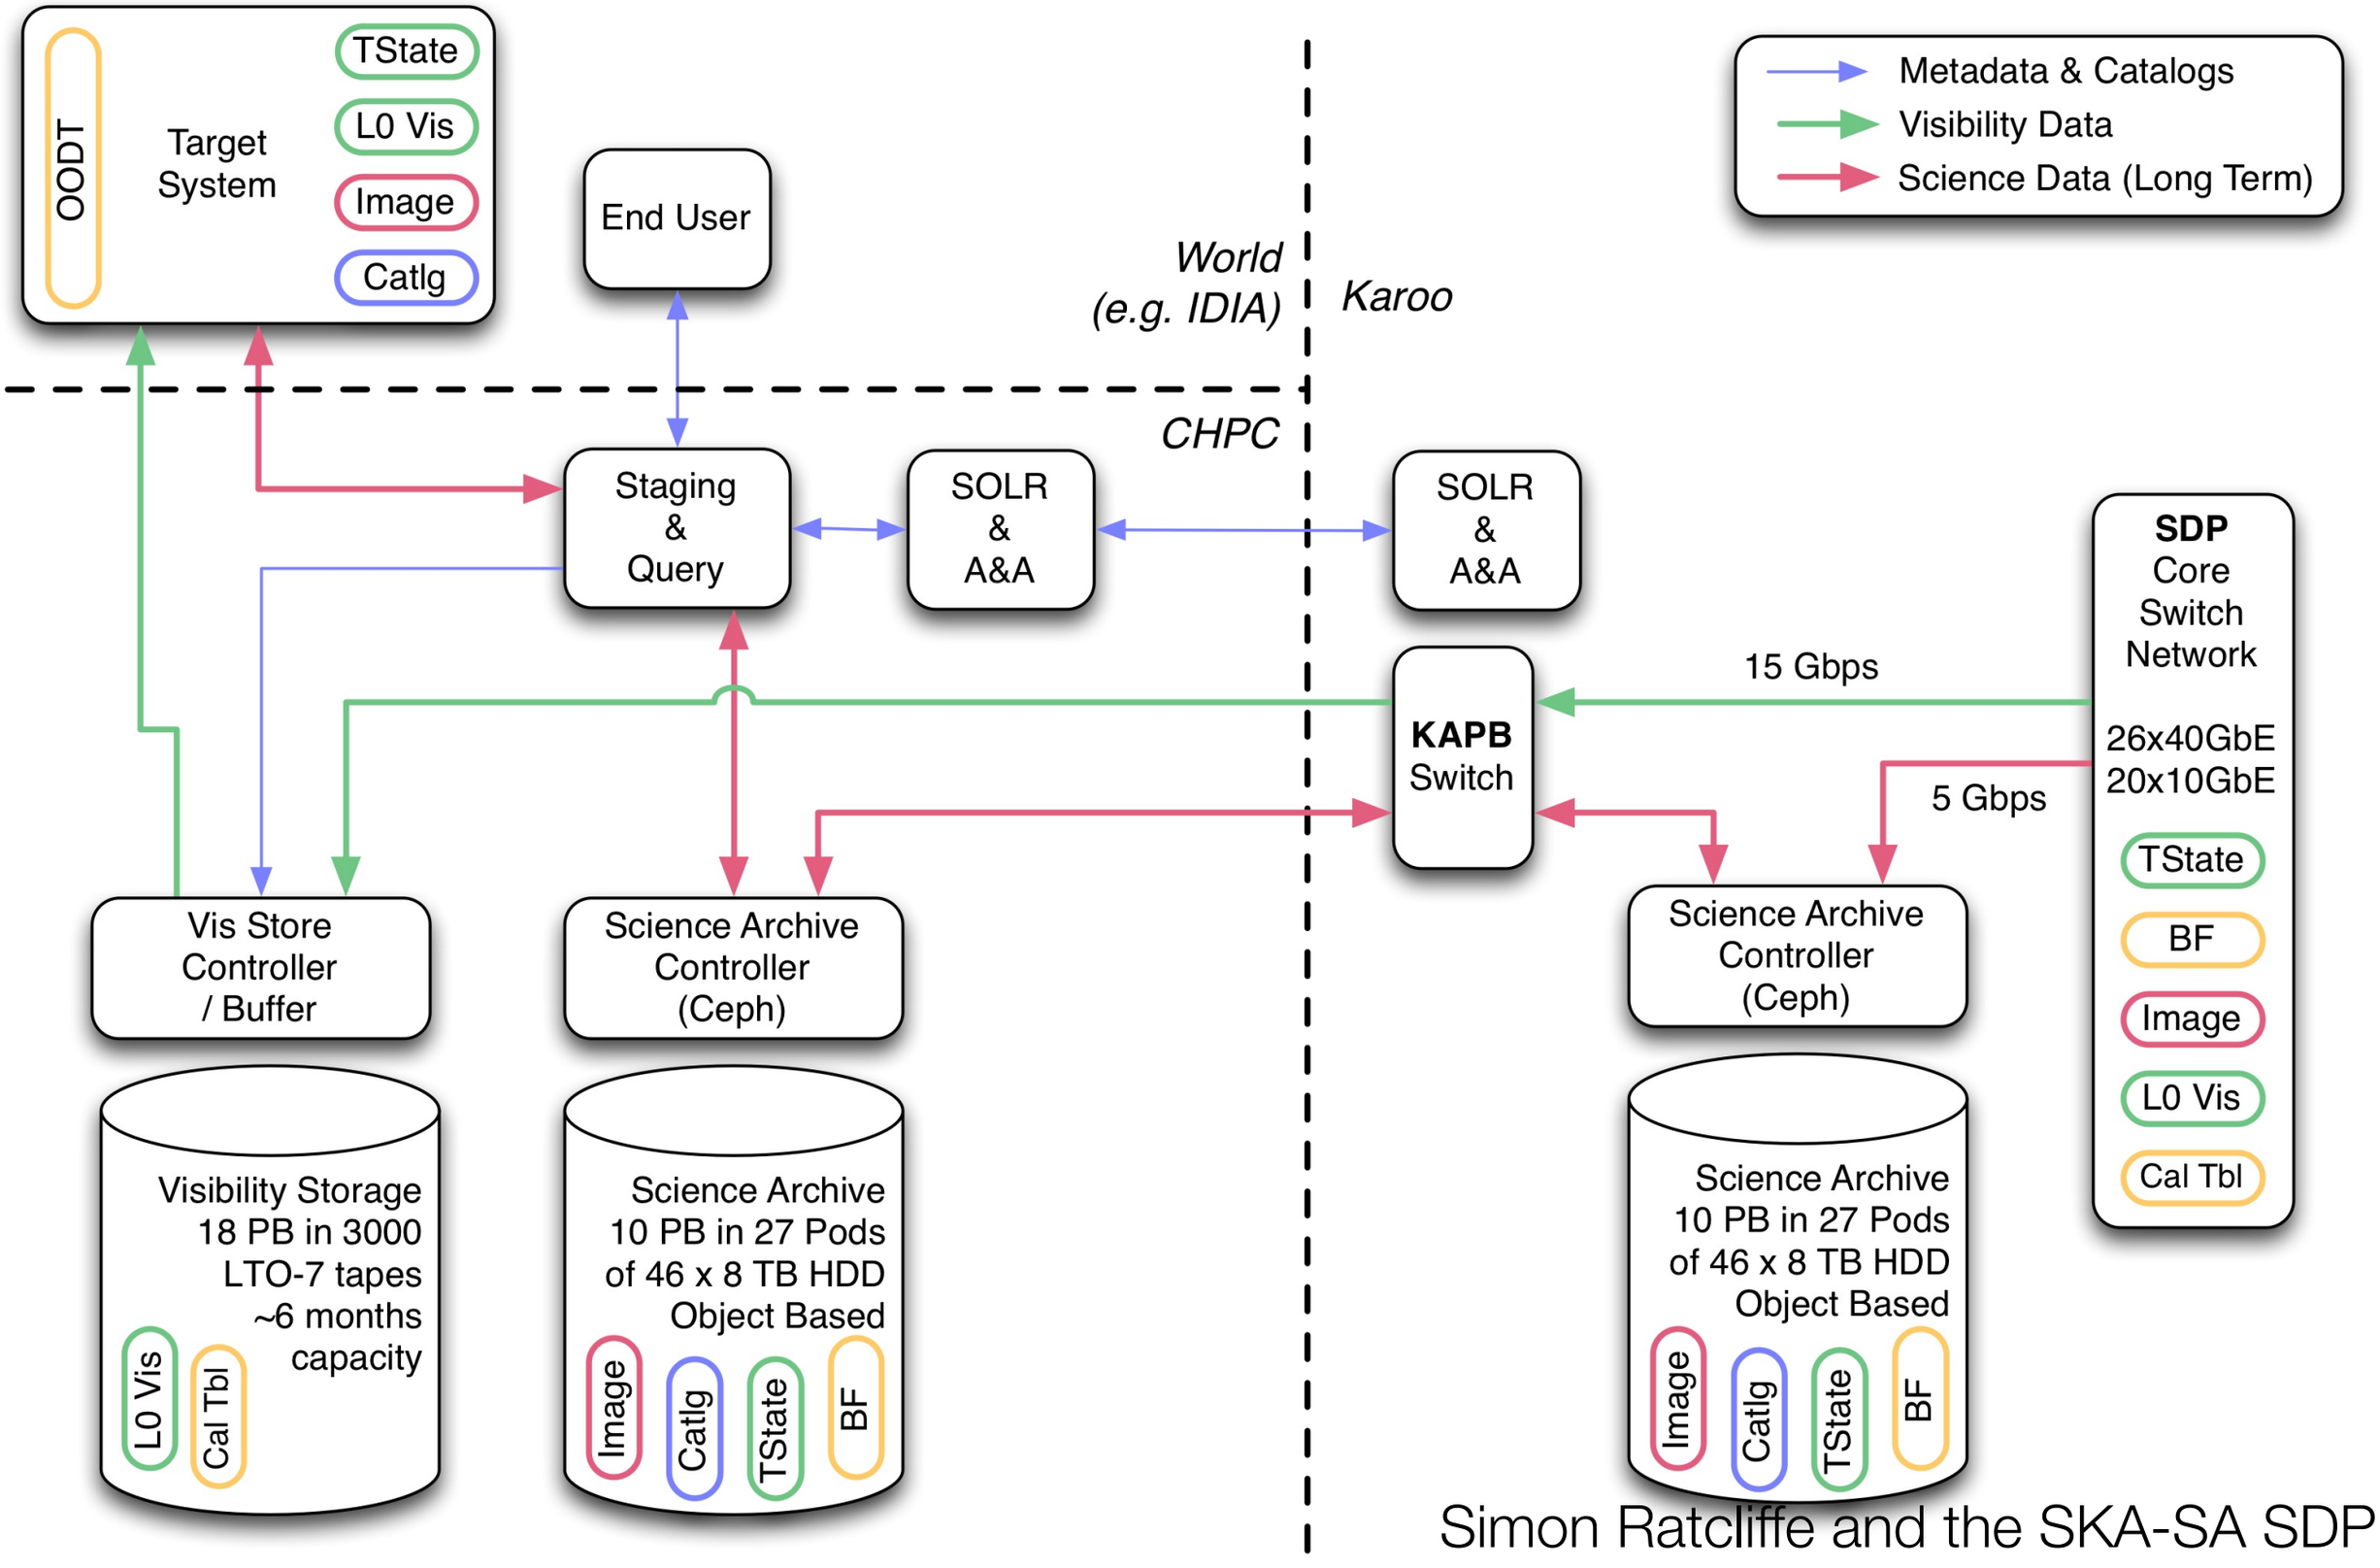

Simon Ratcliffe and the SKA-SA SDP

MeerKAT Data Products

- Simon Ratcliffe and SDP
  - Calibrated, flagged, full time and spectral resolution visibility data (medium term storage) + 10x reduced product (indefinite)
  - Full res is 0.5 Hz / 32,768 channels + per vis flags
  - HDF5 native with MSv2 and FITS export (subset)
  - Vis Data + Cal Tables

Science With MeerKAT

- Science surveys need a home for processing, analysis and calibration — beyond SDP.
- Inter-University Institute for Data Intensive Astronomy (IDIA).
- Partners: UCT, UWC, UP, NWU.
- Support and resources for MeerKAT science community.
- Pathfinder for SKA Regional Science Data Centre (partner with SKA-SA, ASTRON and IBM-Dome).
- Federated cloud based data centre.

IDIA Inter-University Institute for Data Intensive Astronomy
UNIVERSITY OF CAPE TOWN IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD
ARC AFRICAN RESEARCH CLOUD
SKA SKA AFRICA

IDIA

- First roll-out of data centre is up:
  - 40 compute nodes.
    - 2.6GHz Xeon processors.
    - 32 cores.
    - 256GB RAM.
    - 4 nodes with 2x NVidia P100 GPUs.
- Combination of POSIX, Block and Object Store.
- 0.5PB Initial storage (to be expanded in 2018)
- 10Gb/s network access to MeerKAT Archive.
- Located at UCT.
- OpenStack, with Singularity containers and Jupyter Hub.


science
& technology
Department:
Science and Technology
REPUBLIC OF SOUTH AFRICA
WEST UNIVERSITY
UNIVERSITI YA BOKONE-BOPHIRIMA
NOORD-WES-UNIVERSITEIT

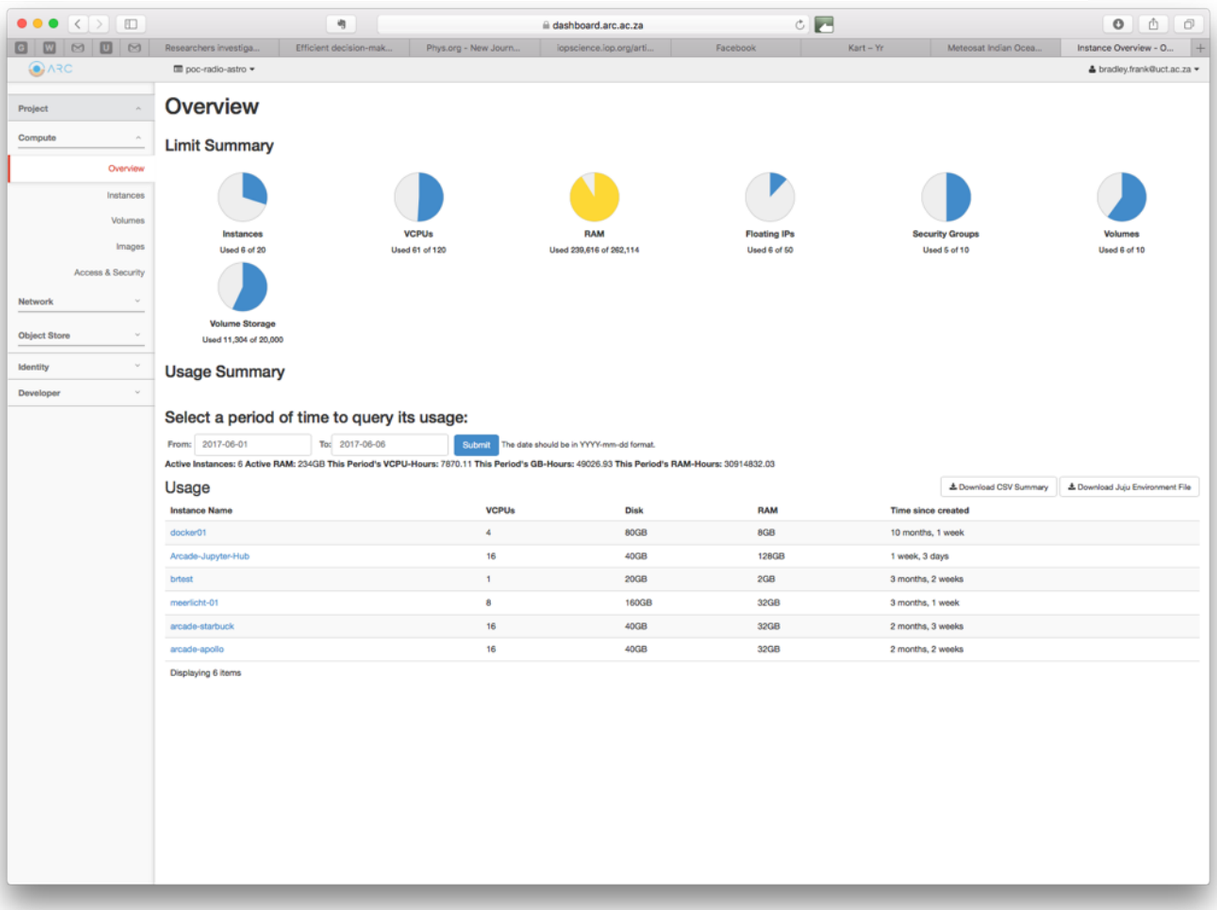
dashboard.arc.ac.za
ARC
poc-radio-astro
bradley.frank@uct.ac.za
Project
Compute
Overview
Instances
Volumes
Images
Access & Security
Network
Object Store
Identity
Developer
Overview
Limit Summary
Instances
Used 6 of 20
VCPUs
Used 61 of 120
RAM
Used 239,616 of 262,114
Floating IPs
Used 6 of 50
Security Groups
Used 5 of 10
Volumes
Used 6 of 10
Volume Storage
Used 11,504 of 20,000
Usage Summary
Select a period of time to query its usage:
From: 2017-06-01
To: 2017-06-06
Submit
The date should be in YYYY-mm-dd format.
Active Instances: 6 Active RAM: 234GB This Period's VCPU-Hours: 7870.11 This Period's GB-Hours: 49026.93 This Period's RAM-Hours: 30914832.03
Usage
Download CSV Summary
Download Juju Environment File
Instance Name
VCPUs
Disk
RAM
Time since created
docker01
4
80GB
8GB
10 months, 1 week
Arcade-Jupyter-Hub
16
40GB
128GB
1 week, 3 days
brtest
1
20GB
2GB
3 months, 2 weeks
meerlicht-01
8
160GB
32GB
3 months, 1 week
arcade-starbuck
16
40GB
32GB
2 months, 3 weeks
arcade-apollo
16
40GB
32GB
2 months, 2 weeks
Displaying 6 items

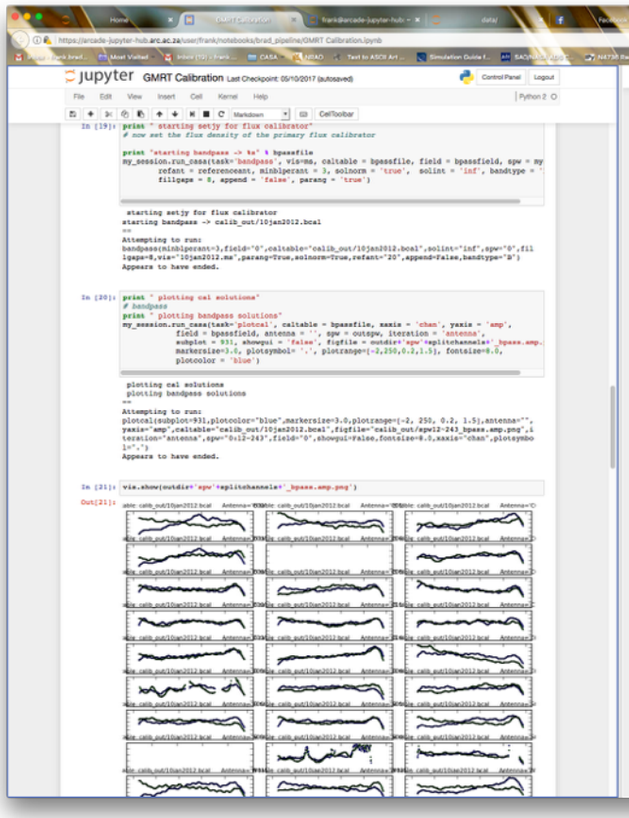
jupyter
GMRT Calibration
Last Checkpoint: 05/10/2017 (autosaved)
Control Panel
Logout
File
Edit
View
Insert
Cell
Kernel
Help
Python 2
CellToolbar
print " starting setjy for flux calibrator"
now set the flux density of the primary flux calibrator
print "starting bandpass -> %s" % bpassfile
my_session.run_casa(task='bandpass', vis=ms, caltable = bpassfile, field = bpassfield, spw = my
refant = referenceant, minblperant = 3, solnorm = 'true', solint = 'inf', bandtype = '
fillgaps = 0, append = 'false', parang = 'true')
starting setjy for flux calibrator
starting bandpass -> calib_out/10jan2012.bcal
Attempting to run:
bandpass(minblperant=3,field="0",caltable="calib_out/10jan2012.bcal",solint="inf",spw="0",fil
lgaps=0,vis="10jan2012.ms",parang=True,solnorm=True,refant="20",append=False,bandtype="B")
Appears to have ended.
print " plotting cal solutions"
bandpass
print " plotting bandpass solutions"
my_session.run_casa(task='plotcal', caltable = bpassfile, xaxis = 'chan', yaxis = 'amp',
field = bpassfield, antenna = '', spw = outspw, iteration = 'antenna',
subplot = 931, showgui = 'false', figfile = outdir+'spw'+splitchannels+'_bpass.amp.
markersize=3.0, plotsymbol='.', plotrange=[-2,250,0.2,1.5], fontsize=8.0,
plotcolor = 'blue')
plotting cal solutions
plotting bandpass solutions
Attempting to run:
plotcal(subplot=931,plotcolor="blue",markersize=3.0,plotrange=[-2, 250, 0.2, 1.5],antenna="",
yaxis="amp",caltable="calib_out/10jan2012.bcal",figfile="calib_out/spw12-243_bpass.amp.png",i
teration="antenna",spw="0:12-243",field="0",showgui=False,fontsize=8.0,xaxis="chan",plotsymbo
l='.')
Appears to have ended.
vis.show(outdir+'spw'+splitchannels+'_bpass.amp.png')

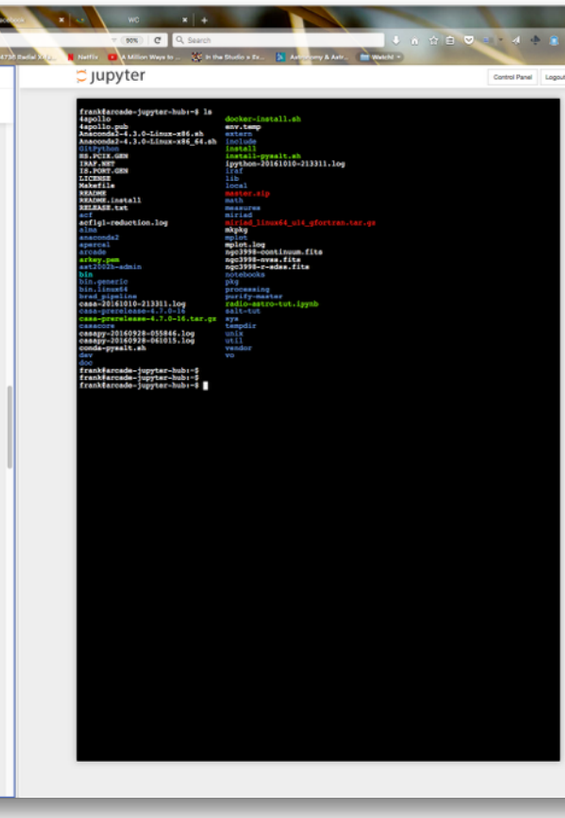
jupyter
Control Panel
Logout
frank@arcade-jupyter-hub:~$ ls
4apollo
4apollo.pub
Anaconda2-4.3.0-Linux-x86.sh
Anaconda2-4.3.0-Linux-x86_64.sh
CASAPython
BB.PCIX.GEN
IRAF.NET
IB.PORT.GEN
LICENSE
Makefile
README
README.install
RELEASE.txt
anf
aoflgl-reduction.log
atms
anaconda2
aperal
arcade
arkey.pem
art2012b-admin
bin
bin.generic
bin.linux64
brad_pipeline
casa-20161010-213311.log
casa-prerelease-4.7.0-18
casa-prerelease-4.7.0-16.tar.gz
casacore
casapy-20160928-055846.log
casapy-20160928-061015.log
conda-pyasit.sh
dev
doc
docker-install.sh
env.temp
extern
include
install
install-pyasit.sh
ipython-20161010-213311.log
lib64
lib
local
master.zip
math
measures
miriad
miriad_linux64_u14_gfortran.tar.gz
mkpkg
mplot
mplot.log
ngc3998-continuum.fits
ngc3998-cvsa.fits
ngc3998-z-edax.fits
notebooks
pkg
processing
purify-master
radio-astro-tut.ipynb
salt-tut
sys
tempdir
unix
util
vendor
vo
frank@arcade-jupyter-hub:~$
frank@arcade-jupyter-hub:~$
frank@arcade-jupyter-hub:~$

Typical Data Flow

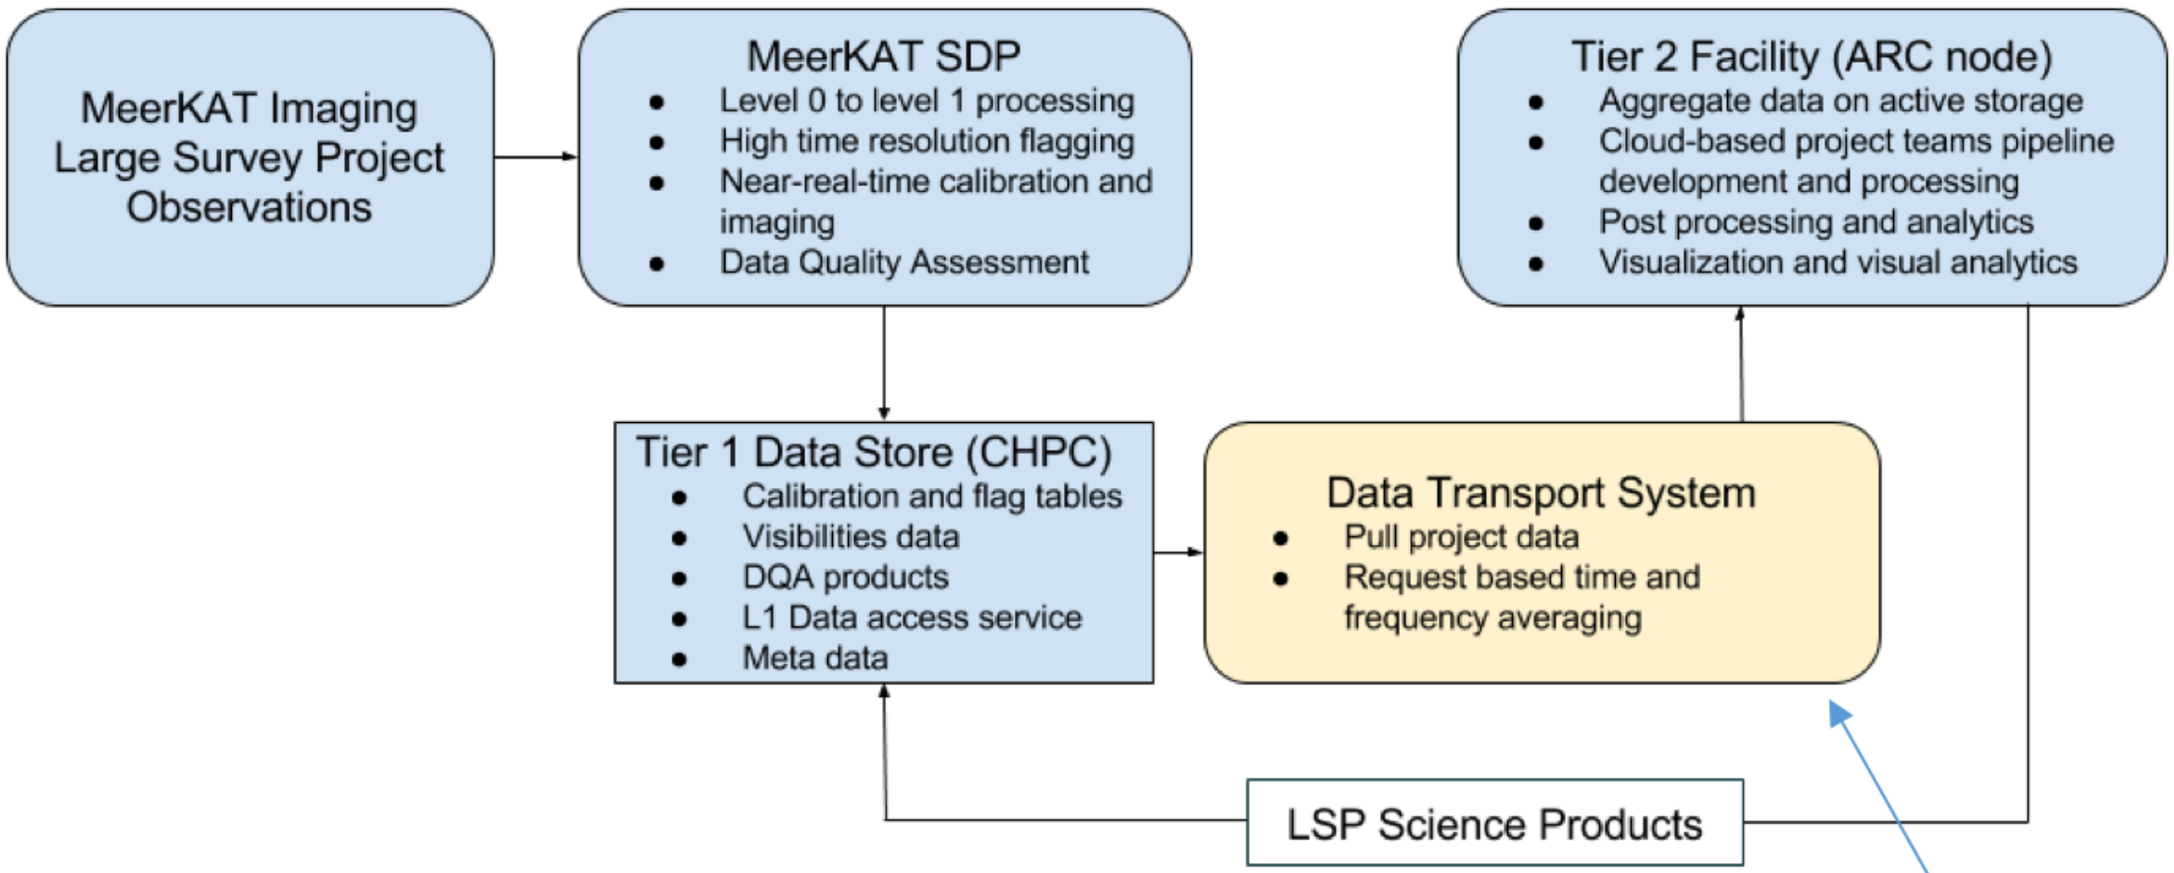

Being developed by IDIA and SKA-SA

IDIA

- MoU with NRAO to co-develop CASAPY functionality.
- SAP: industrial partner.
- ASTRON/IBM-Dome/SKA-SA/IDIA Pathfinder Science Regional Data Centre.
- Many MLSPs have co-Is (or PIs) at IDIA institutions.
- Initial proposal for project based resources (*virtual* hardware & people).
- Transition to Tier 2 Data Intensive Research Cloud (supported by South African DST).

Data Transport Project
Based on SKA DELIV

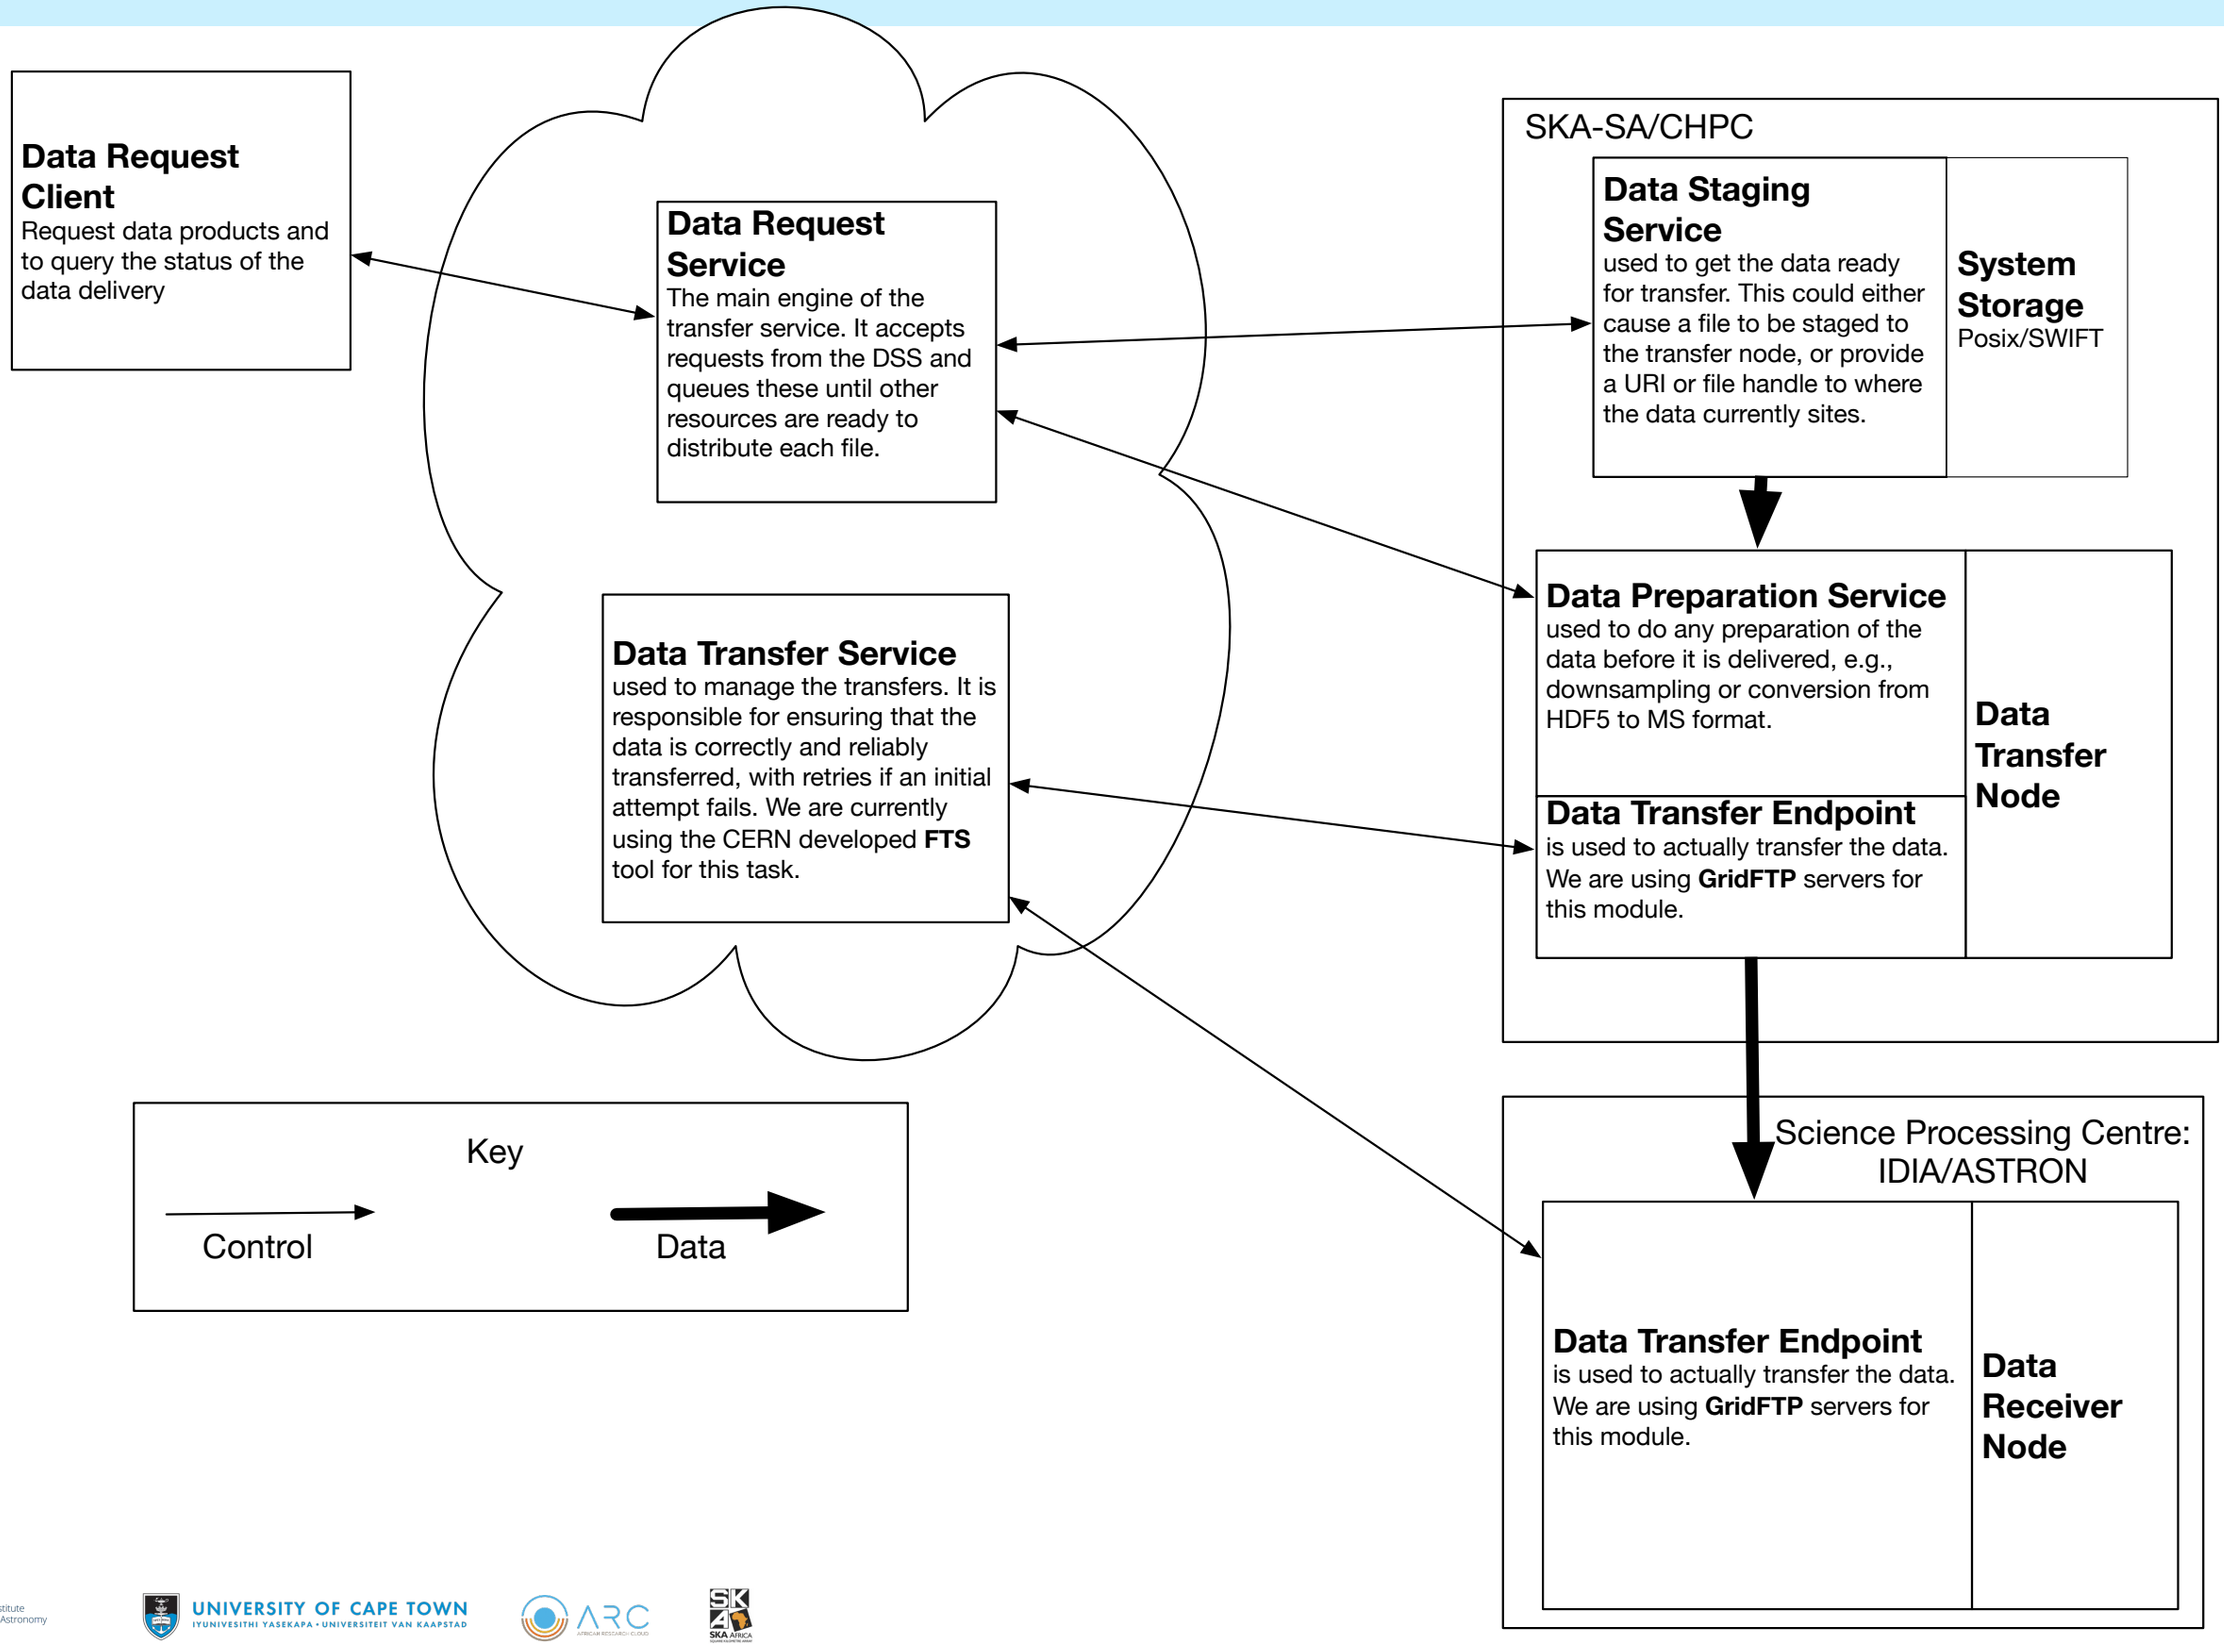

RSDC & Data Transport.

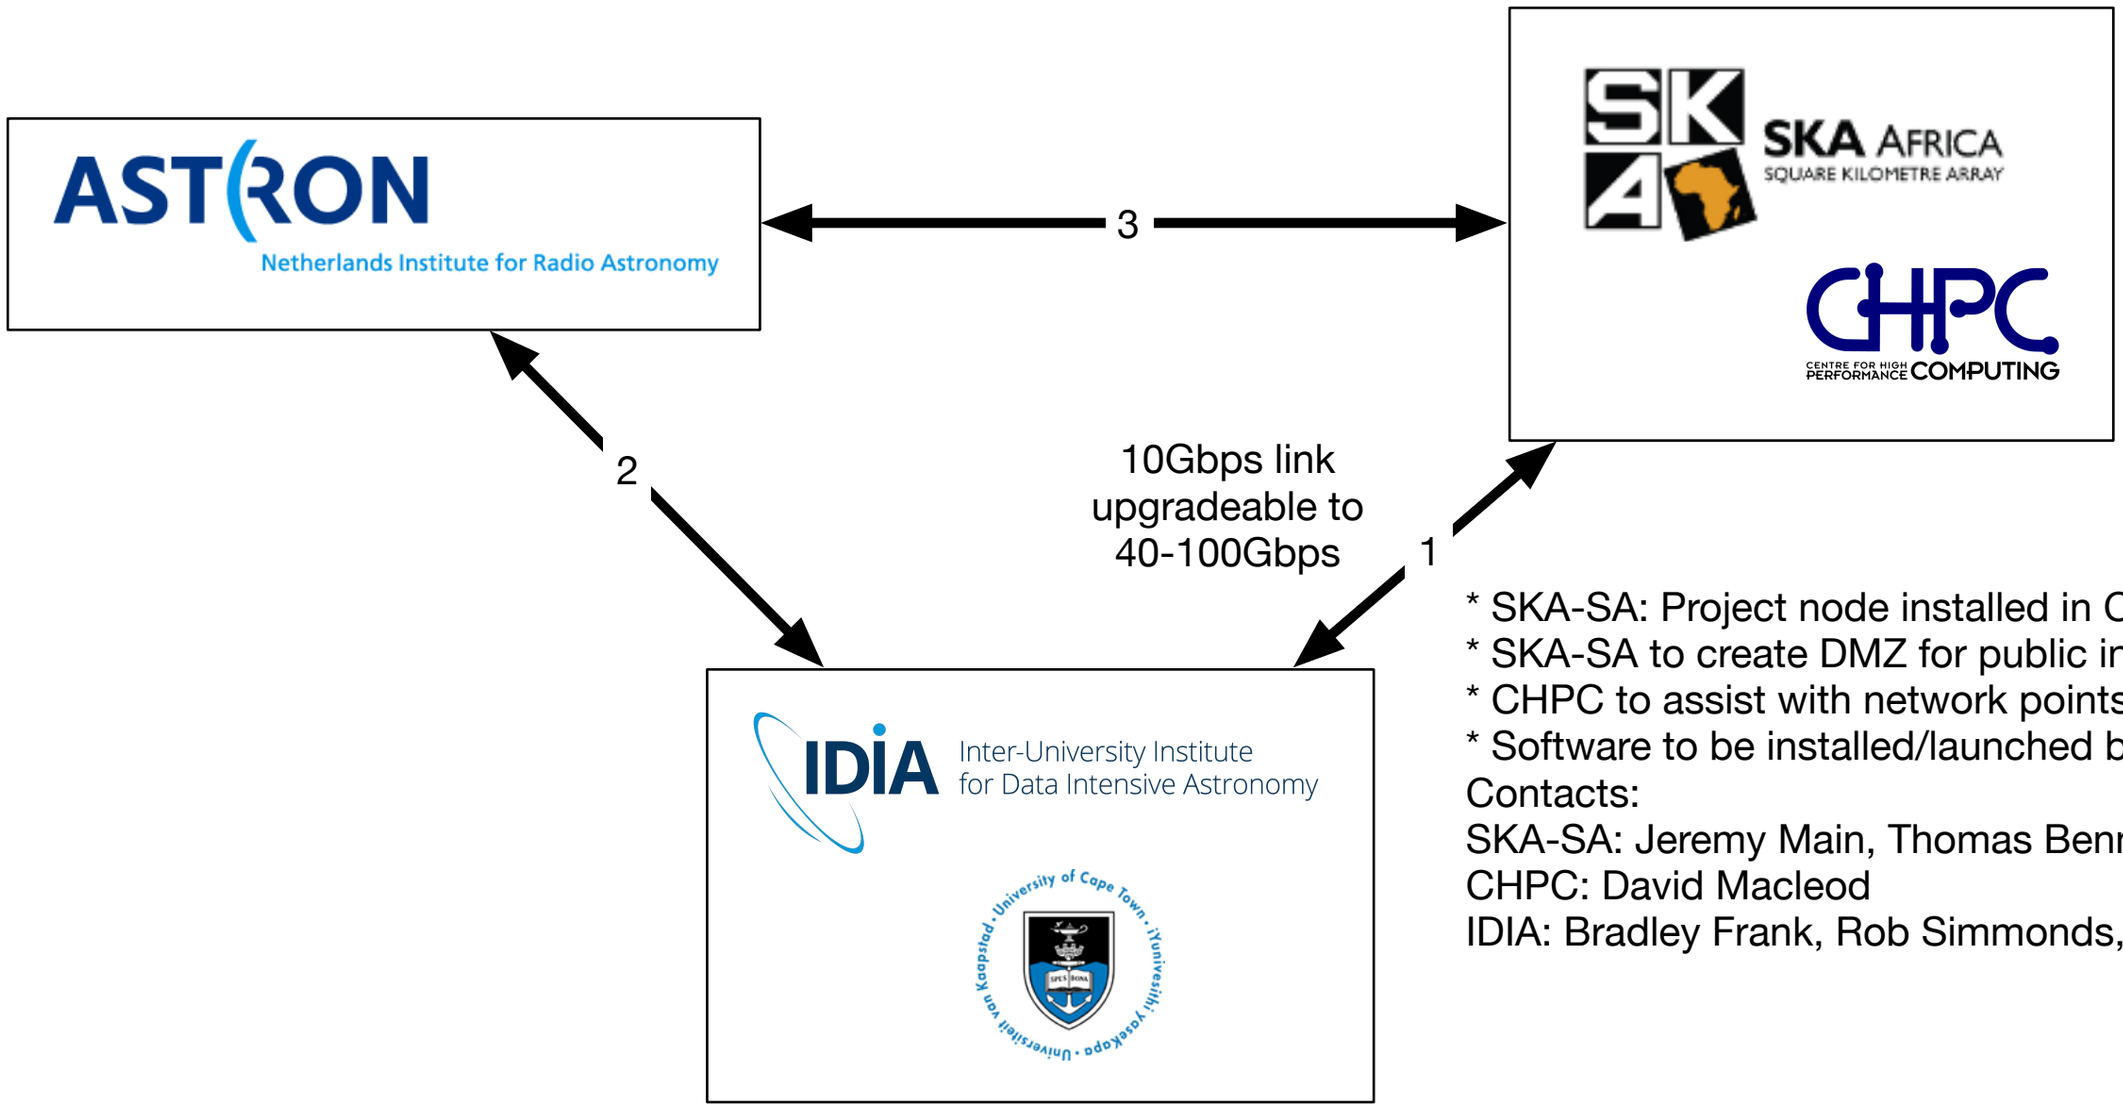

* SKA-SA: Project node installed in CHPC.
* SKA-SA to create DMZ for public interface.
* CHPC to assist with network points.
* Software to be installed/launched by IDIA.

Contacts:
SKA-SA: Jeremy Main, Thomas Bennett, Simon Ratcliffe
CHPC: David Macleod
IDIA: Bradley Frank, Rob Simmonds, David Aikema

Request for Access/Resources

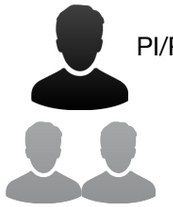

PI/PUs

Support

Support

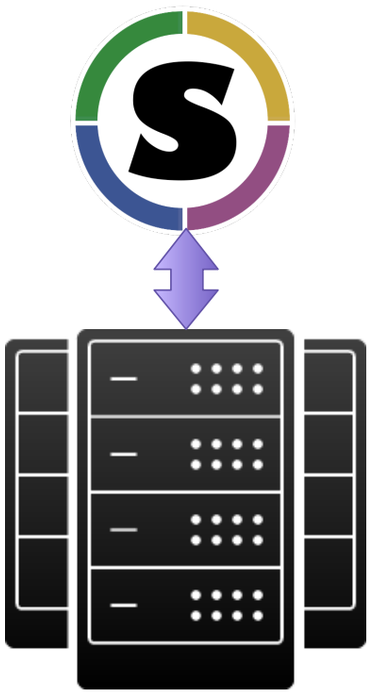

Fat Node

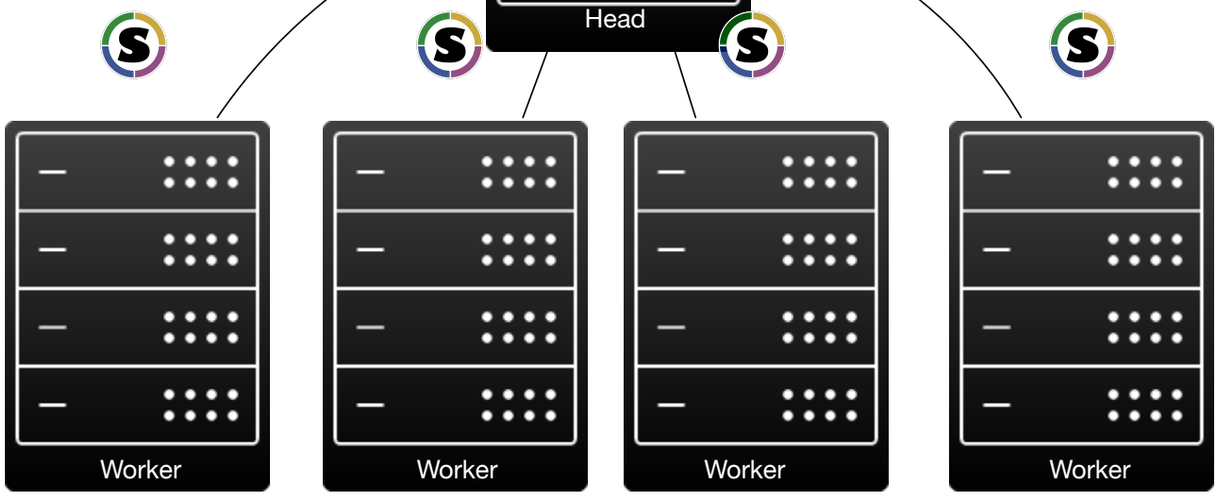

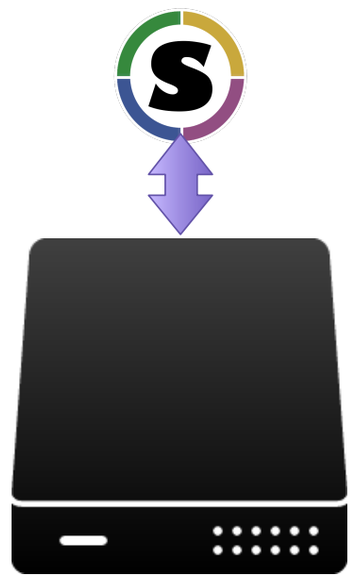

Bare Metal

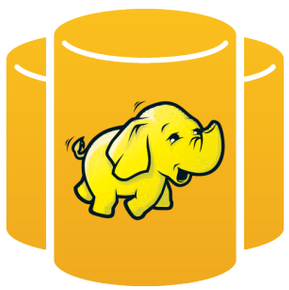

Hadoop/Distributed Storage

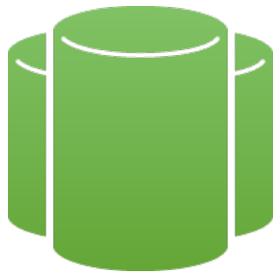

Posix

Data Processing

- Currently — prototyping: using single nodes.
- Near future — deployment: Pipeline converted into a graph using Common Workflow Language.
- Resources allocated per *node*: Mesos/SLURM/

Dagman…

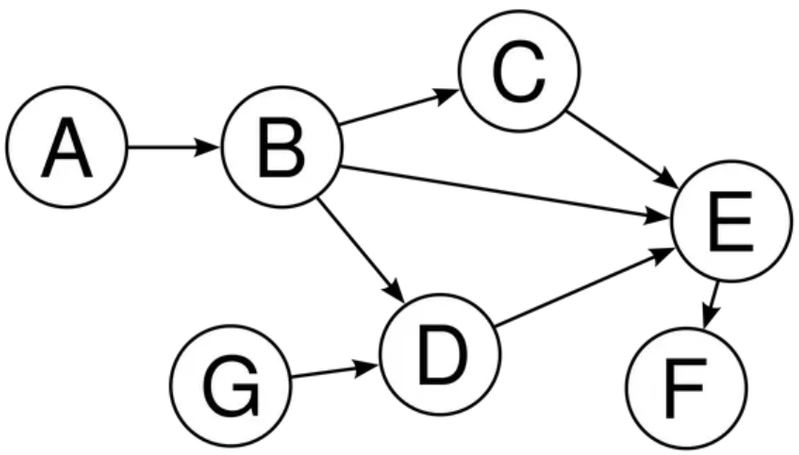

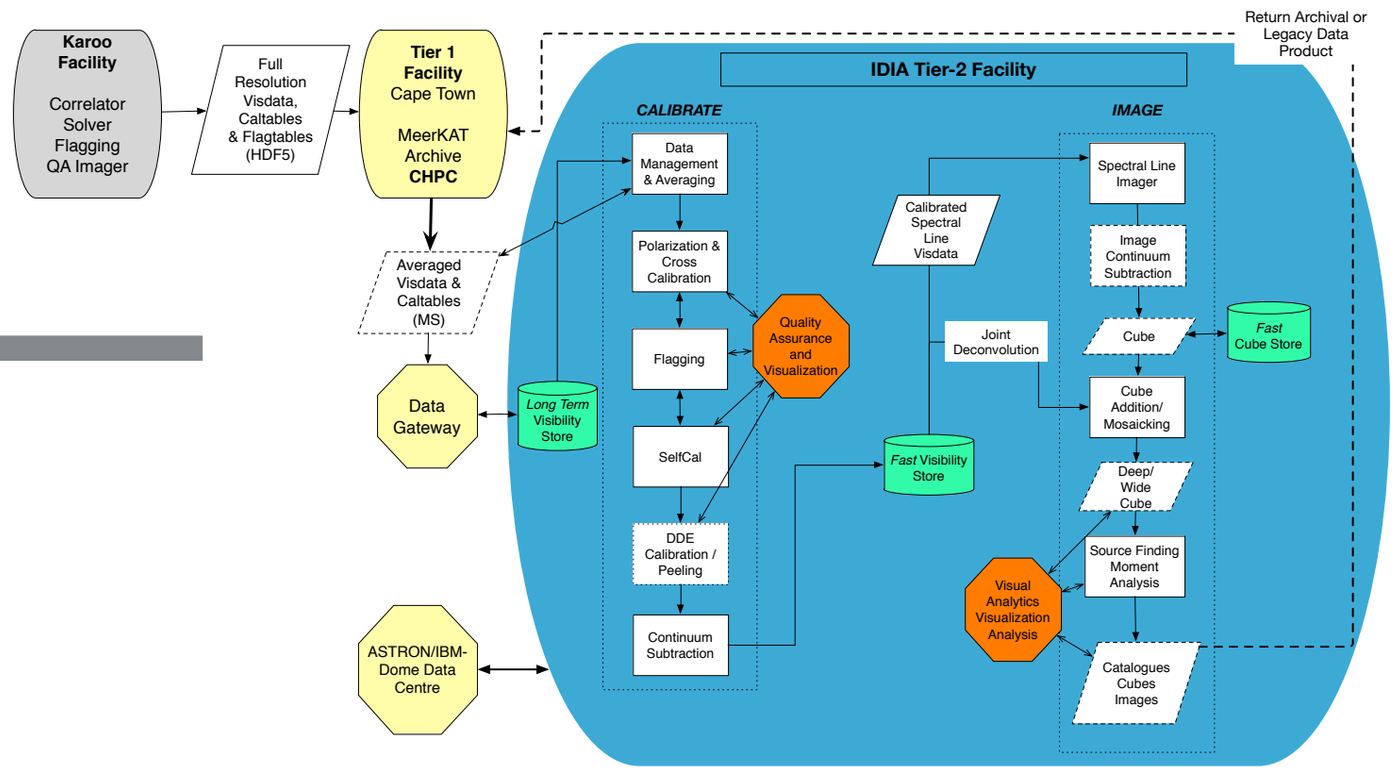

IDIA Inter-University Institute for Data Intensive Astronomy
UNIVERSITY OF CAPE TOWN IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD
ARC AFRICAN RESEARCH CLOUD
SKA AFRICA

Conclusions

- MeerKAT AR1.5 data is available to MLSPs for pipeline development.
- Commissioning/Science-verification with MLSPs to start soon.
- MeerKAT AR3 ready by April 2018.
- Early science data soon thereafter.
- IDIA: data processing, analysis and collaboration.
  - Pathfinder for the SKA Regional Science Data Centre.