Big Data Management in MeerKAT+

The MeerKAT Telescope

- 64 13.5m dishes with an offset gregorian configuration;
- each dish has 3 receivers at different bands;
- 875 MHz instantaneous bandwidth;
- minimum baseline: 29 m
- maximum baseline: 8 km

The MeerKAT telescope

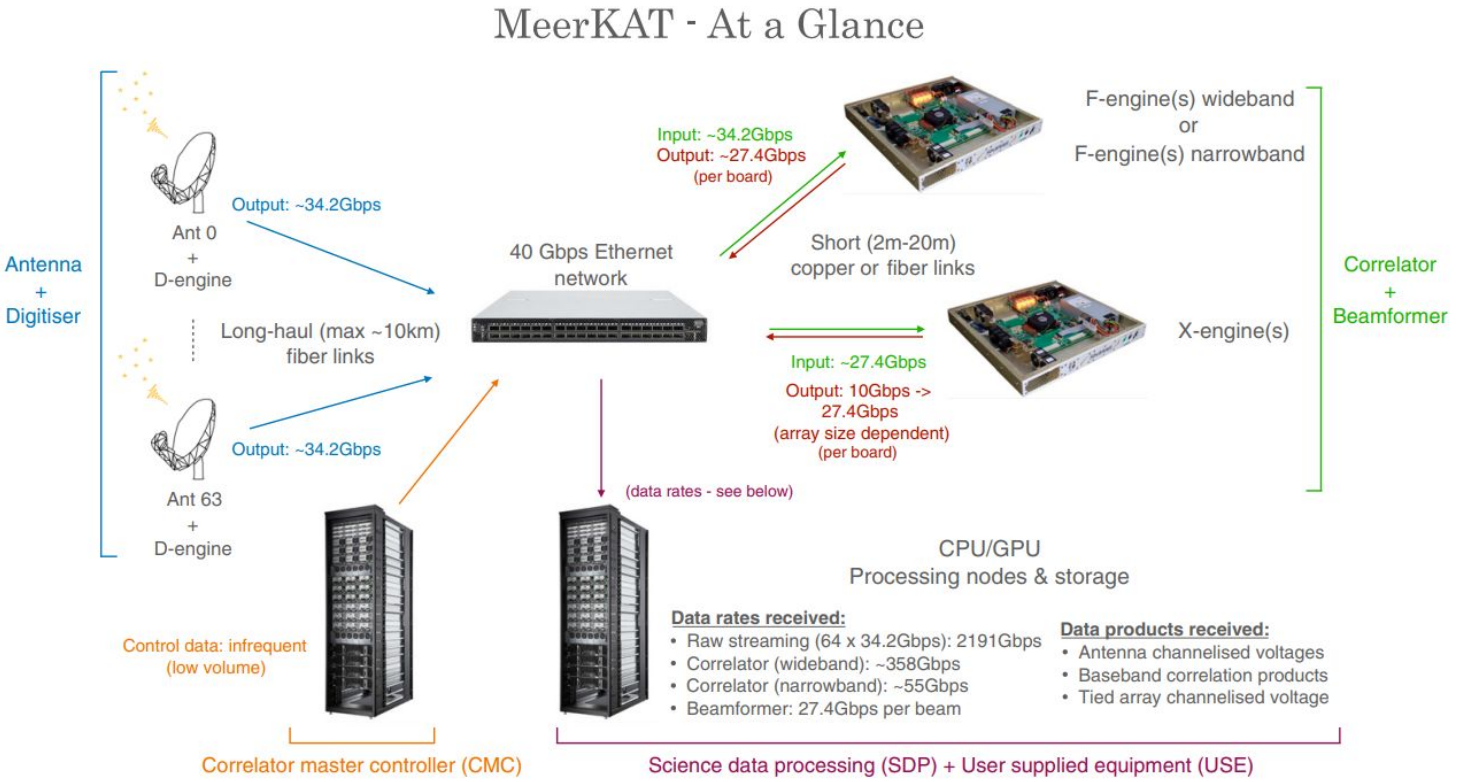

van der Byl et al., 2022

The MeerKAT telescope

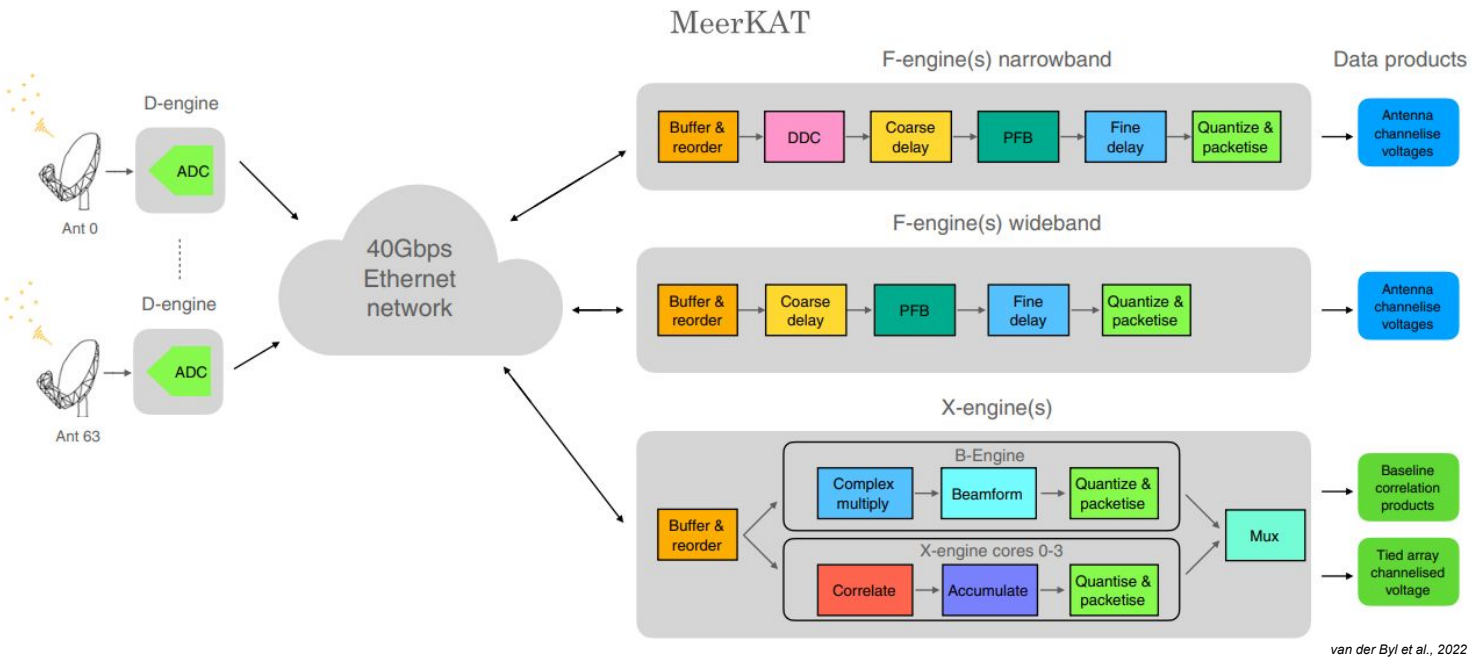

van der Byl et al., 2022

Measurement set

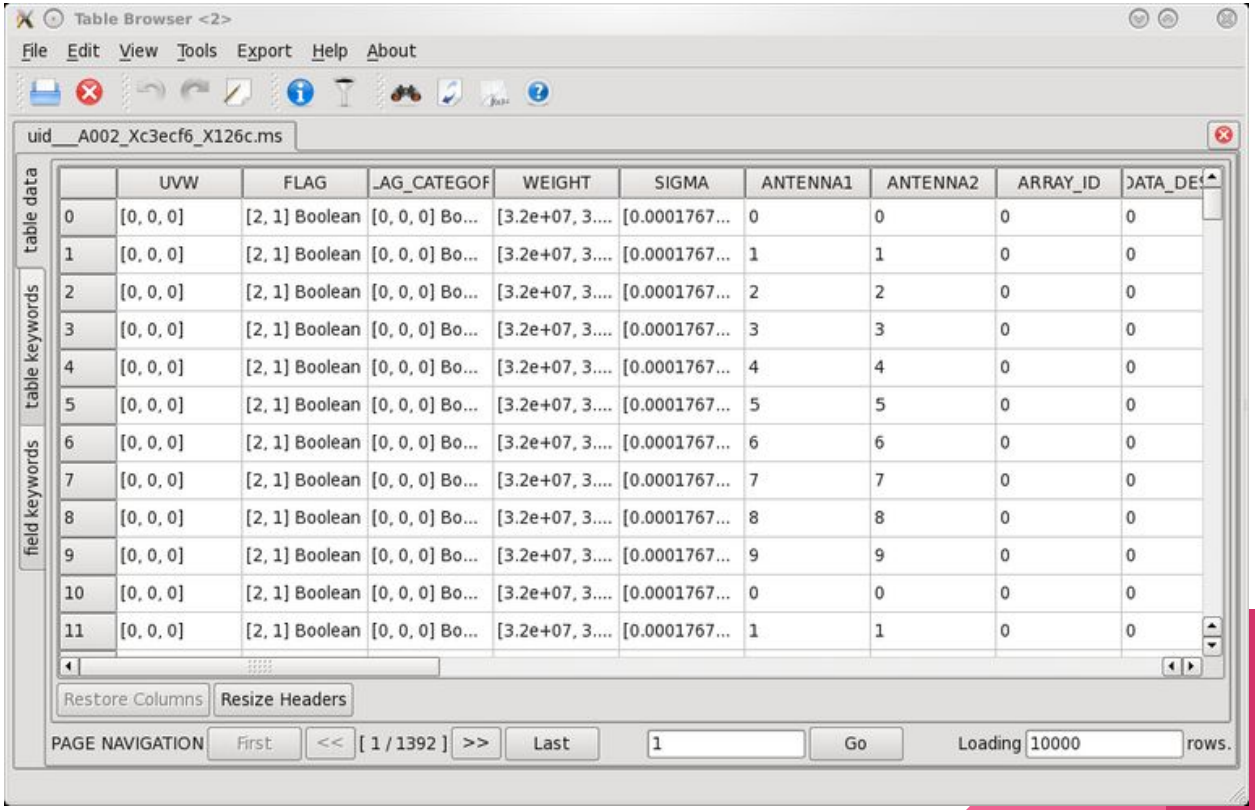

	UVW	FLAG	_AG_CATEGOF	WEIGHT	SIGMA	ANTENNA1	ANTENNA2	ARRAY_ID	)ATA_DE
0	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	0	0	0	0
1	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	1	1	0	0
2	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	2	2	0	0
3	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	3	3	0	0
4	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	4	4	0	0
5	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	5	5	0	0
6	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	6	6	0	0
7	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	7	7	0	0
8	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	8	8	0	0
9	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	9	9	0	0
10	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	0	0	0	0
11	[0, 0, 0]	[2, 1] Boolean	[0, 0, 0] Bo...	[3.2e+07, 3....	[0.0001767...	1	1	0	0

Measurement set

How big is one hour of MeerKAT observation?

- N*(N-1)/2 (2016) baselines times
- 4 correlations (xx,yy,xy,yx) times
- integration intervals (one every 2-4-8 seconds) (1800-900-450) times
- number of channels (4k-32k)

0.2 TB (4k-8 seconds) - 4 TB (32k-4 seconds)

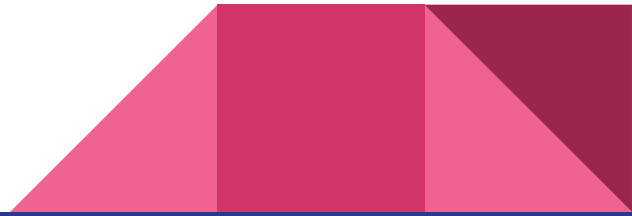

Science Data Processing

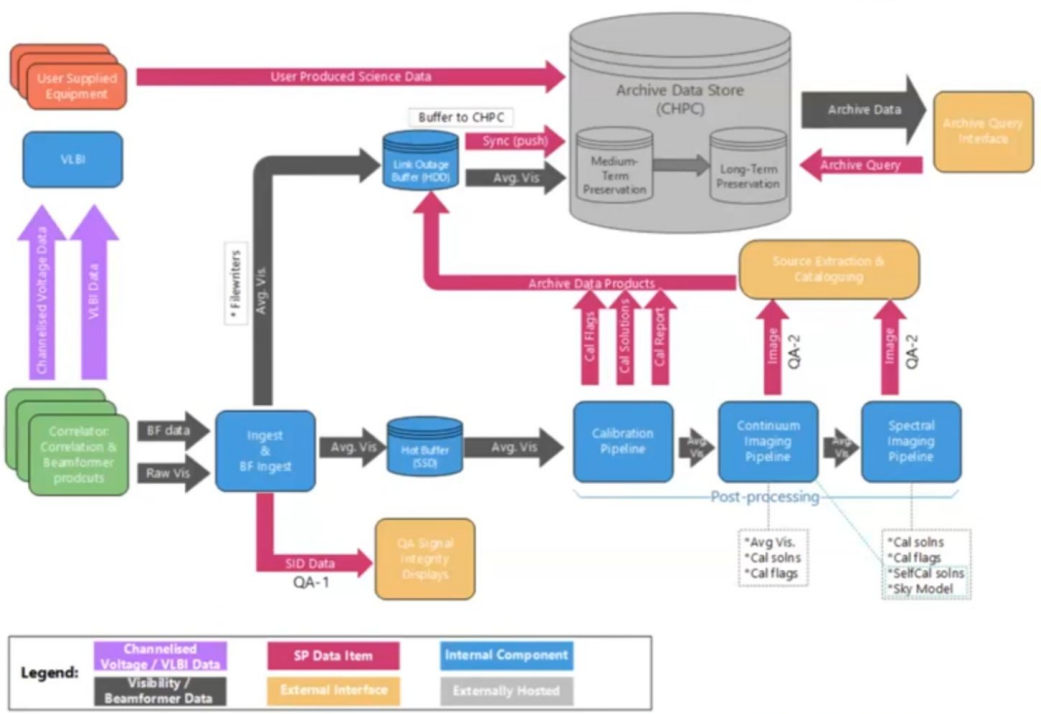

The ilifu cluster

Partition	Node names	Default CPUs	Max CPUs	Default Memory (GiB)	Max Memory (GiB)	Default wall-time	Max wall-time
Main	compute-[002-021]	1	32	3	232	3 hours	14 days
Main	compute-[101-105]	1	48	3	232	3 hours	14 days
Main	compute-[201-260]	1	32	3	251	3 hours	14 days
HighMem	highmem-[001-002]	1	32	15	503	3 hours	14 days
HighMem	highmem-003	1	96	15	1508	3 hours	14 days
GPU	gpu-[001-004]	1	32	7	232	3 hours	14 days
GPU	gpu-005	1	24	7	232	3 hours	14 days
GPU	gpu-006	1	48	7	354	3 hours	14 days
GPU	gpu-007	1	48	7	354	3 hours	14 days
Devel	compute-001	1	32	-	-	3 hours	12 hours

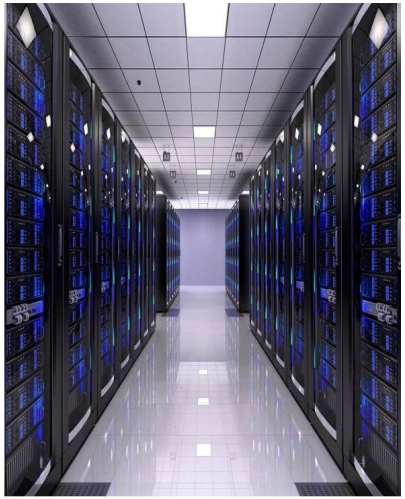

MeerKAT data reduction pipelines

- Science Data Processing (SDP, from SARAO)
- processMeerKAT (pipeline IDIA)
- oxKAT (used in the Galactic Center images)
- caracal (Rhodes University)

Compute-intensive, several days of processing for a single night, using several nodes on a cluster due to the large number of FFT to execute.

The Galactic Center with MeerKAT

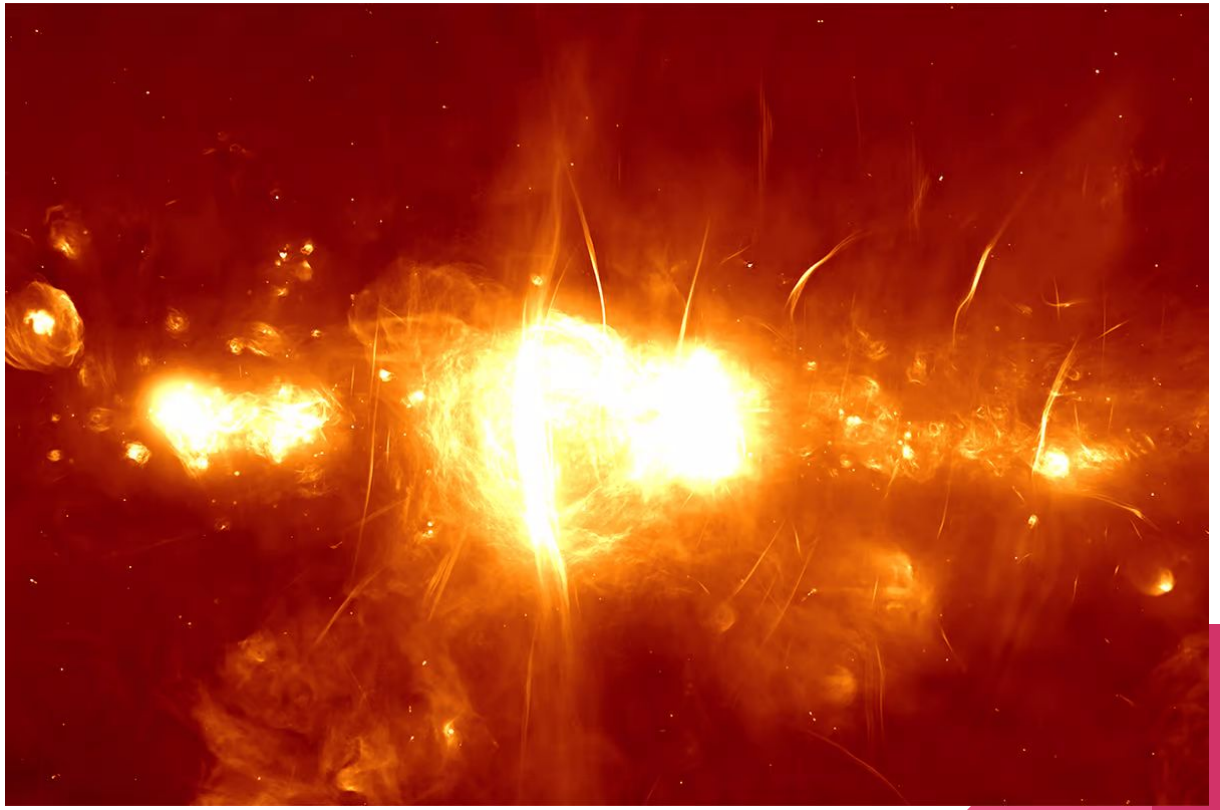

The MeerKAT extension (MK+)

- Started in 2019 as a joint SARAO/MPIfR project
- INAF joined in 2020
- Goal: to develop exciting novel capabilities for MeerKAT, in the context of eventual contributions to SKA-Mid
- Enormous project, drawing on significant resources in: systems engineering, infrastructure, dish structures, receivers, vacuum/cryogenics, fiber optic networks, digitizers, correlator, science data processor, control and monitoring, commissioning…

https://www.meerkatplus.tel/documents/

MeerKAT+

- 14 new antennas, of ‘SKA design’ (each ~1.35x more sensitive than MeerKAT’s)
  - Maximum baselines ~17 km (~twice as much as MeerKAT)
  - Populated with L- and S-band receivers (not UHF)
    ‘L’ are SKA Band-2 receivers (0.95–1.76 GHz, compared to 0.9–1.67 GHz for MeerKAT L)
    Only half of the S band (1.75–3.5 GHz) processed at once
- Backend to process all 78 antennas simultaneously
  - New independent GPU-based correlator for all 78
  - Initial wideband correlator mode: 8k channels
  - Upgraded SDP to ingest and store visibilities of this mode
  - Current SKARAB-based (any) 64-antenna correlator with USE will remain fully functional

https://www.meerkatplus.tel/documents/

MeerKAT+

MeerKAT+ configuration

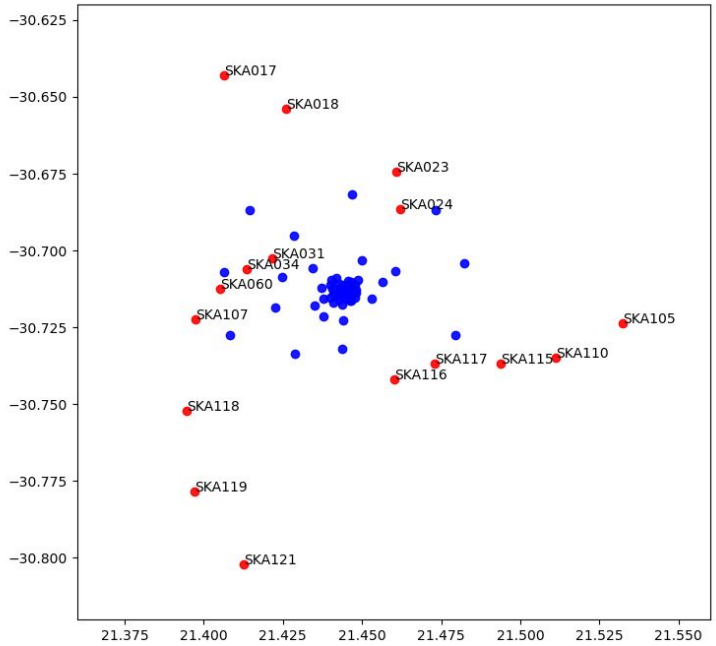

https://www.meerkatplus.tel/documents/

MeerKAT+ receivers

MK+ Receiver Systems SKA Band 2 and MeerKAT digitizer:

- L-Band (0.90 – 1.67 GHz) vs Band 2 (0.95 – 1.76GHz);
- cryogenic cooled: LNA: 20K, OMT: 70K

MPIfR S-Band

- frequency range 1.75 – 3.50 GHz, digitally reduced to 0.875 GHz bandwidth to cater for correlator;
- cryogenic cooled: LNA and OMT <20K.

Key science projects

Current White Paper (WP), released on May 2021, presents ideas for 3 legacy science projects. INAF did not participate/contribute to the discussion that led to the drafting. While broad in their science goals, these three surveys can be summarised as:

(1) a MeerKAT HI galaxy evolution survey;
(2) a full-Stokes S-band southern sky survey ($\delta < -400$)
(3) a cosmological HI survey at L-band, including the intensity-mapping technique to probe large-scale structure

All three are aligned with SKA-MID key science programmes and contribute to SKA Key Science Projects survey design in the future

New challenges arise

The CASPER FPGA based correlators currently operating for the MeerKAT telescope can not work on more than 64 antennas.

To address it they are considering changing the correlators, upgrading to a GPU-based system

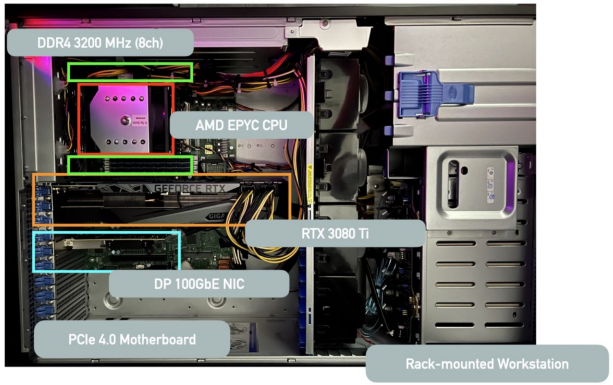

New challenges arise

The higher resolution generates an effect (called smearing) that can be eliminated only by lowering the time interval. The 8 second default time interval will not be feasible. These, together with the higher number of antennas, will generate a larger stream of data, that will be partly addressed by using an 8k channels digitalization.

Data Storage system spans multiple Technologies and Geographical sites:

1) KDRA (KAROO)
   a) CEPH
      i) HDDs
      ii) SSDs

CHPC (Cape Town Site)

- CEPH
  - Current: ~21PiB
    - Custom Made Hardware (Peralex)
      - HDDs
      - SSDs
  - Next gen: > 9.7PiB
    - Supermicro
      - HDD
    - Dell
      - SSD
- IBM Tape Library
  - Raw capacity ~21PiB
  - Usable ~7PiB

New challenges arise

The presence of different dish designs goes against one of the postulate at the basis of almost every data reduction algorithm. There are two algorithms that almost fit our needs (mosaic and awprojection, part of the CASA package) but the first is not suited for wide field observations, while the second was not thought for these kind of heterogeneous array.

To address this issue new algorithms must be developed.

Original project roadmap

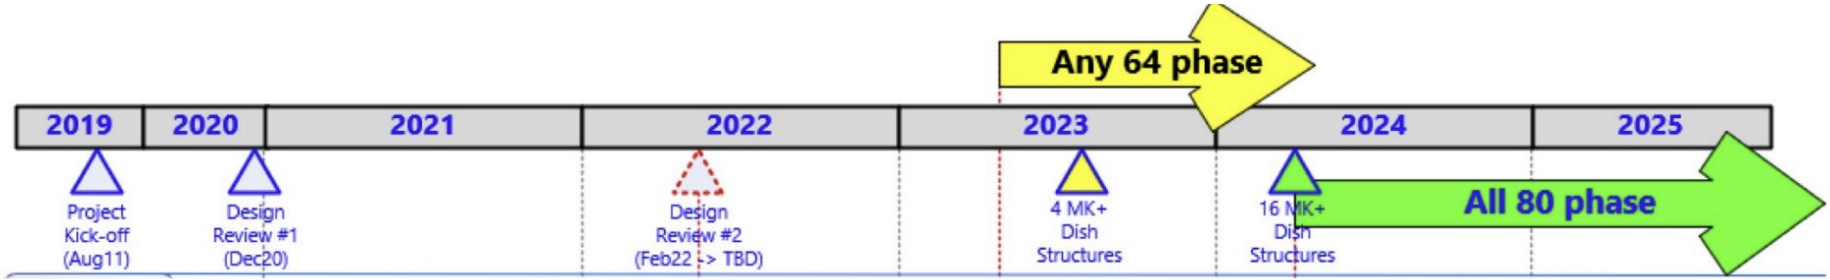

At the moment just one antenna is operative and is actually being tested by the engineers. Commissioning will begin early next year


Thank you!