Entre développement durable et coûts : le projet de radio-télescopes SKA au carrefour de l'efficacité et de la frugalité

Frugalias workshop presentation

Shan Mignot

2024-10-04

Opening statements

- Mission statement: “The SKAO’s mission is to build and operate cutting-edge radio telescopes to transform our understanding of the Universe, and deliver benefits to society through global collaboration and innovation.”

- Sustainability statement: “Sustainability is a foundational value at the SKAO, underpinning all other activities across the Observatory.”

- CO_2 impact of astronomy infrastructures: "We find that worldwide active astronomical research infrastructures currently have a carbon footprint of 20.3 ± 3.3 $MtCO_2$ equivalent (CO_2e) and an annual emission of 1,169 ± 249 $ktCO_2e$/yr”, “research infrastructures make the single largest contribution to the carbon footprint of an astronomer” (Jürgen Knödlseder et al., Nature, 2022).

Mid antenna and layout (courtesy SKAO)

Introduction

Mid telescope

- Frequency range: 350 MHz-15.4 GHz
- 197 dishes, including MeerKAT (red dots)
- Dish geometry: 15 m parabolic reflector
  13.5 m for MeerKAT
- Losberg in the South African Karoo region
- Distribution
  - core within ~1km
  - 3 spiral arms
    - up to 150 km baselines
    - logarithmic distribution

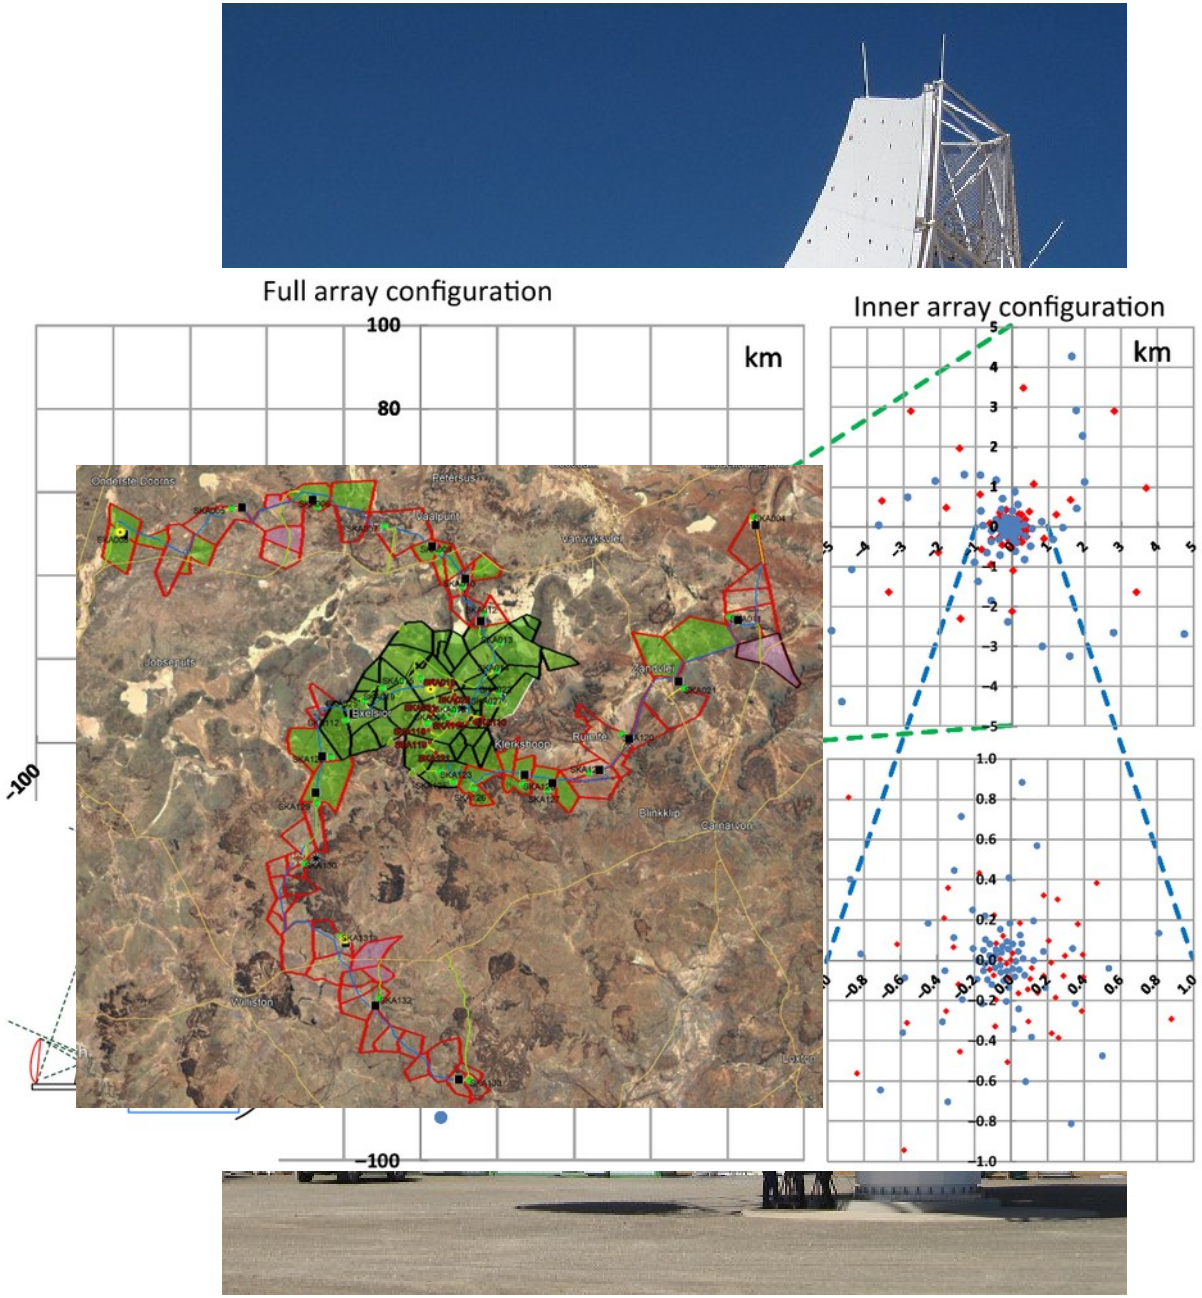

Mid antenna and layout (courtesy SKAO)

Low Telescope

- Frequency range: 50-350 MHz
- 131 072 log-periodic dipole antennas
- Aperture array telescope: 256 antennas per stations, 512 stations
- Boolardy, Western Australia
- Distribution
  - randomised locations
  - core within 4 km (224 stations)
  - 3 spiral arms (288 stations)
    - up to 74km baselines
    - logarithmic distribution

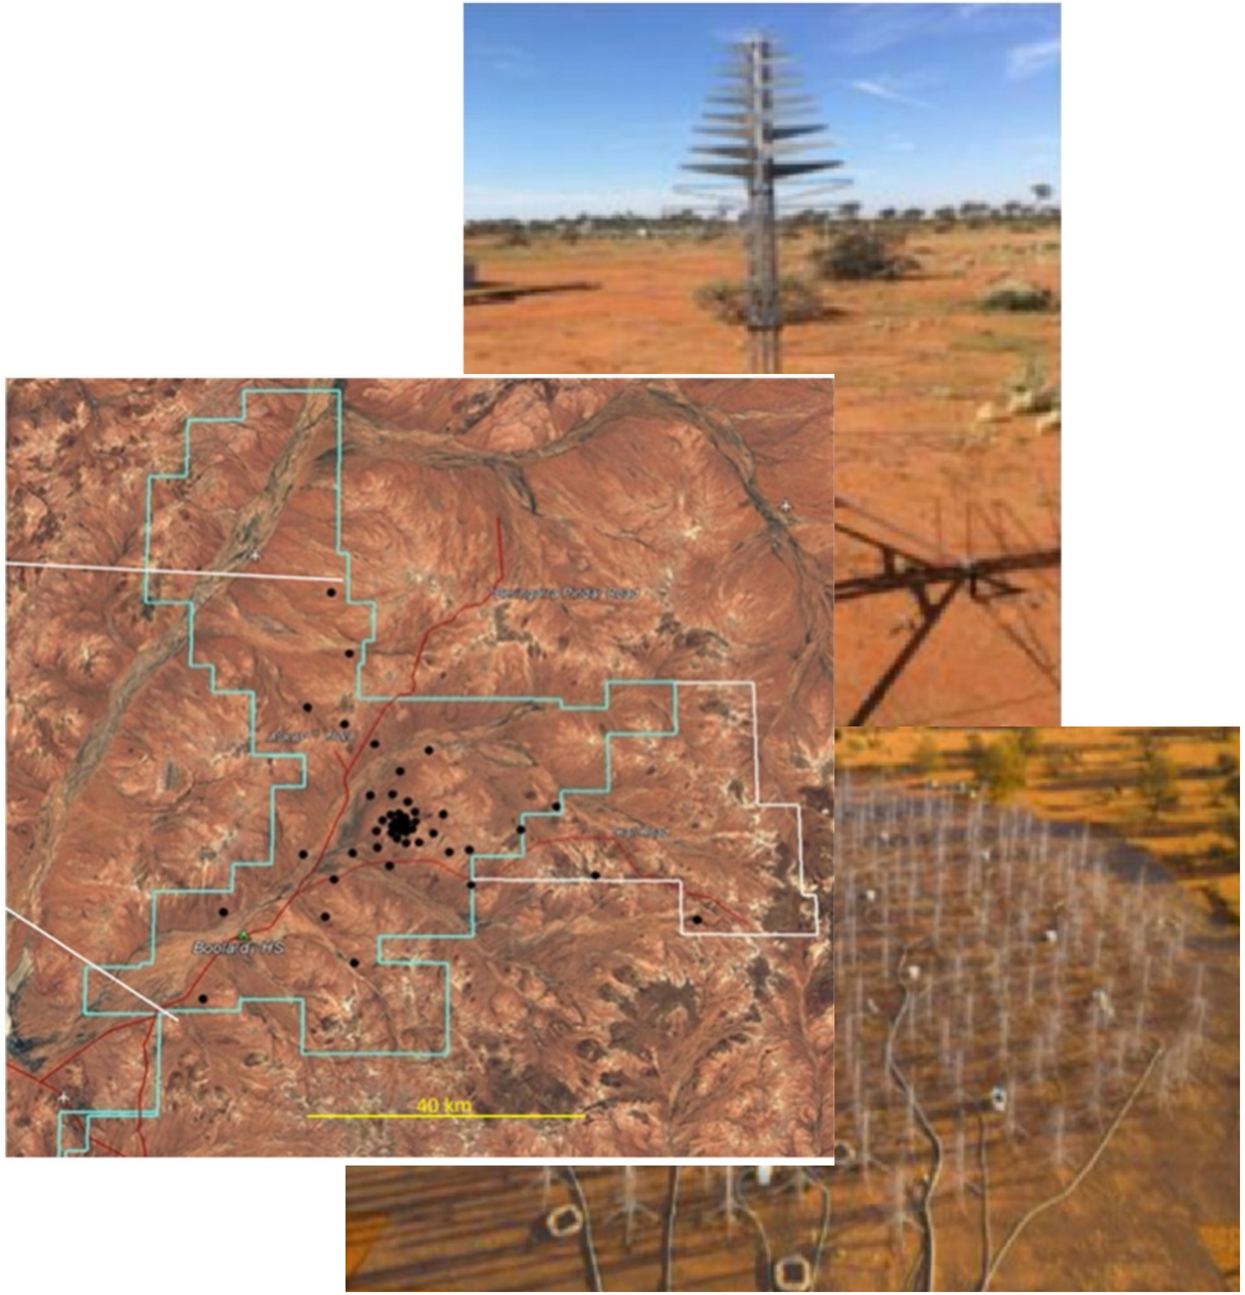

Low antenna and layout (courtesy SKAO)

Construction timeline

- Critical Design Review: 2019-2020
- Start of construction: July 2021
- Progressive deployment of antennas
  - AA0.5: minimal array for early de-risking
  - AA1: de-risking via comparison to existing telescopes
  - AA2: already a full size radio-telescope
    - first science with SKA
    - fully functional system which should scale to SKA1
  - AA*: scaled AA2 offered to the community

Schedule Jan 2024	SKA-Low	SKA-Mid
Start of Construction (T0)	1st July 2021	
Array Assembly 0.5 finish (AA0.5) SKA-Low = 4-station array SKA-Mid = 4-dish array	Nov 2024	May 2025
Array Assembly 1 finish (AA1) SKA-Low = 18-station array SKA-Mid = 8-dish array	Nov 2025	April 2026
Array Assembly 2 finish (AA2) SKA-Low = 64-station array SKA-Mid = 64-dish array	Oct 2026	March 2027
Array Assembly 3* finish (AA*) SKA-Low = 307-station array SKA-Mid = 144-dish array	Jan 2028	Dec 2027
Array Assembly 4 finish (AA4) SKA-Low = 512-station array SKA-Mid = 197-dish array	*N/A*	*N/a*
Operations Readiness Review (ORR)	April 2028	April 2028

Position of the imaging problem

Data acquisition

wide field of view	single pixel
wide sensitivity band	
polarisation sensitive	
operate day and night	

- Detector characteristics: interferometric imaging
  - produce spatial x frequency data cubes (x polarisation)
- Early signal digitization: numerical instruments
  - versatile instrument
  - simplified RF design and improved SNR
  - Low (50-350 MHz) and Mid (350 MHz – 15.4 GHz) and sub-nanosecond pulsar timing: G samples/s
- Interferometry
  - for each pair of antenna calculate visibilities by correlating signals
  - per channel: 12 bytes complex visibility per integration time step (> 0.14 s)

Interferometric imaging (I)

- Visibilities are samples in the Fourier plane
  - (u, v) coordinates correspond to antenna layout
  - coverage is increased by the Earth's rotation
  - coverage is incomplete and irregular
- Invert
  - grid visibilities to rely on FFT
  - produces the "dirty image"
    - artifacts related to (u, v) coverage and missing information
    - complex PSF

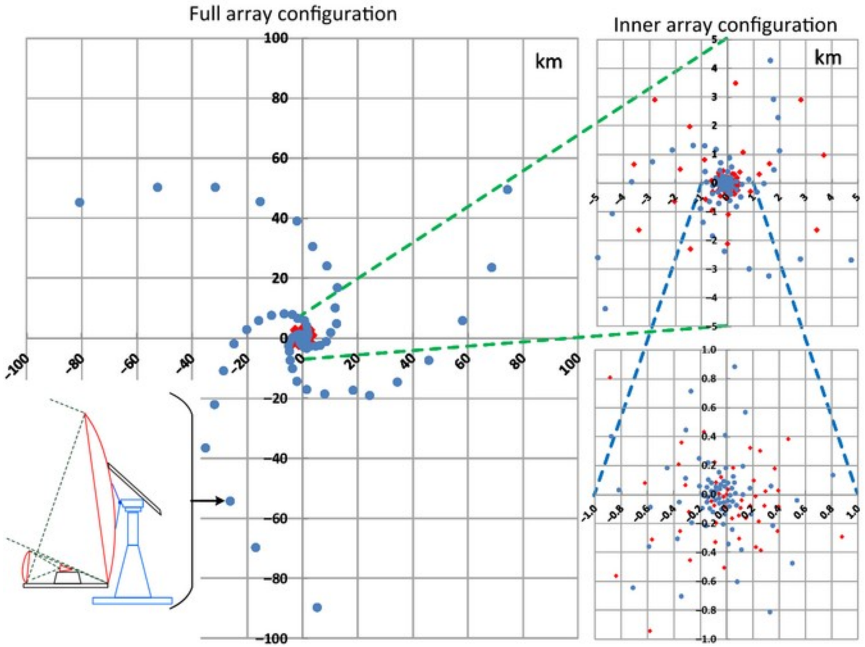

Mid antenna layout (courtesy SKAO)

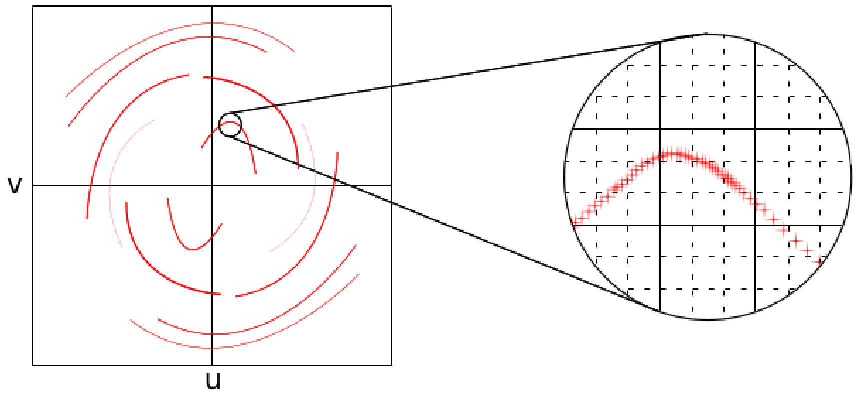

uv coverage resulting Earth rotation (courtesy Nicolas Monnier)

Interferometric imaging (II)

- Cleaning
  - deconvolve to extract sources
  - iterative process: “minor loop”
- Predict visibilities
  - based on source list (or deconvolved image)
  - subtract from input visibilities
  - iterate: “major loop”

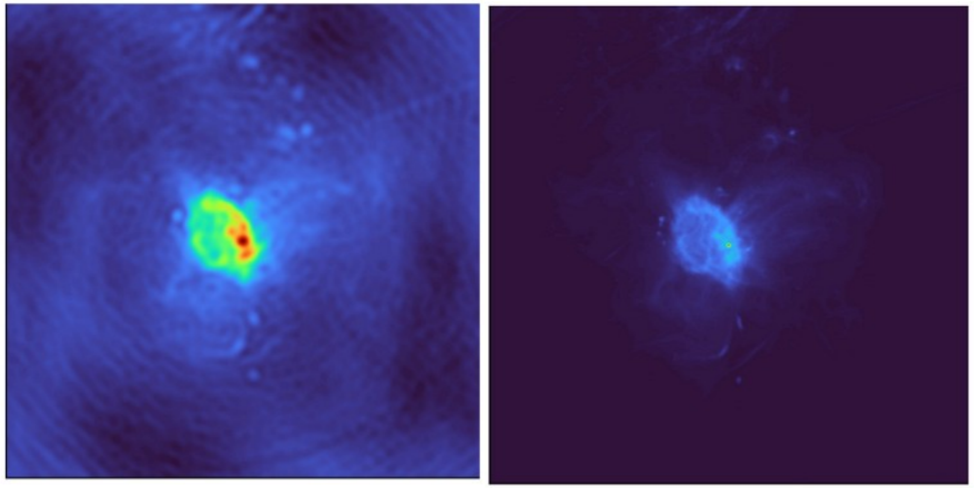

dirty and final images of Sgr A (courtesy Sunrise Wang)

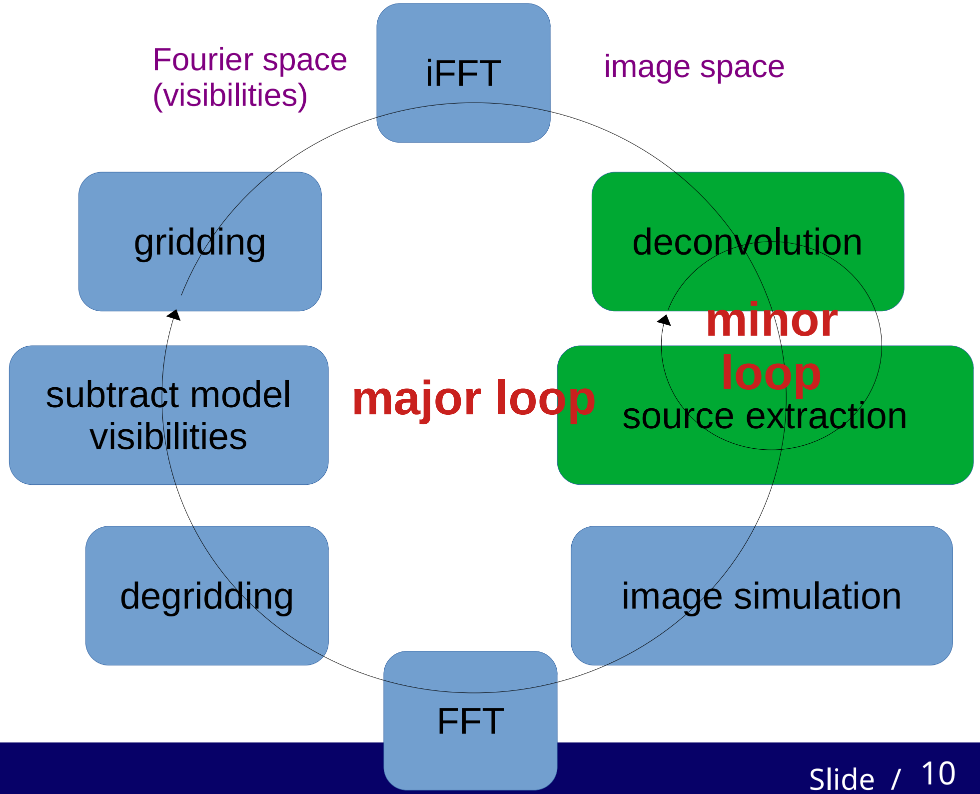

Interferometric imaging (III)

- Complete process involves building the sky and instrument models
  - self-calibration involves an additional calibration loop

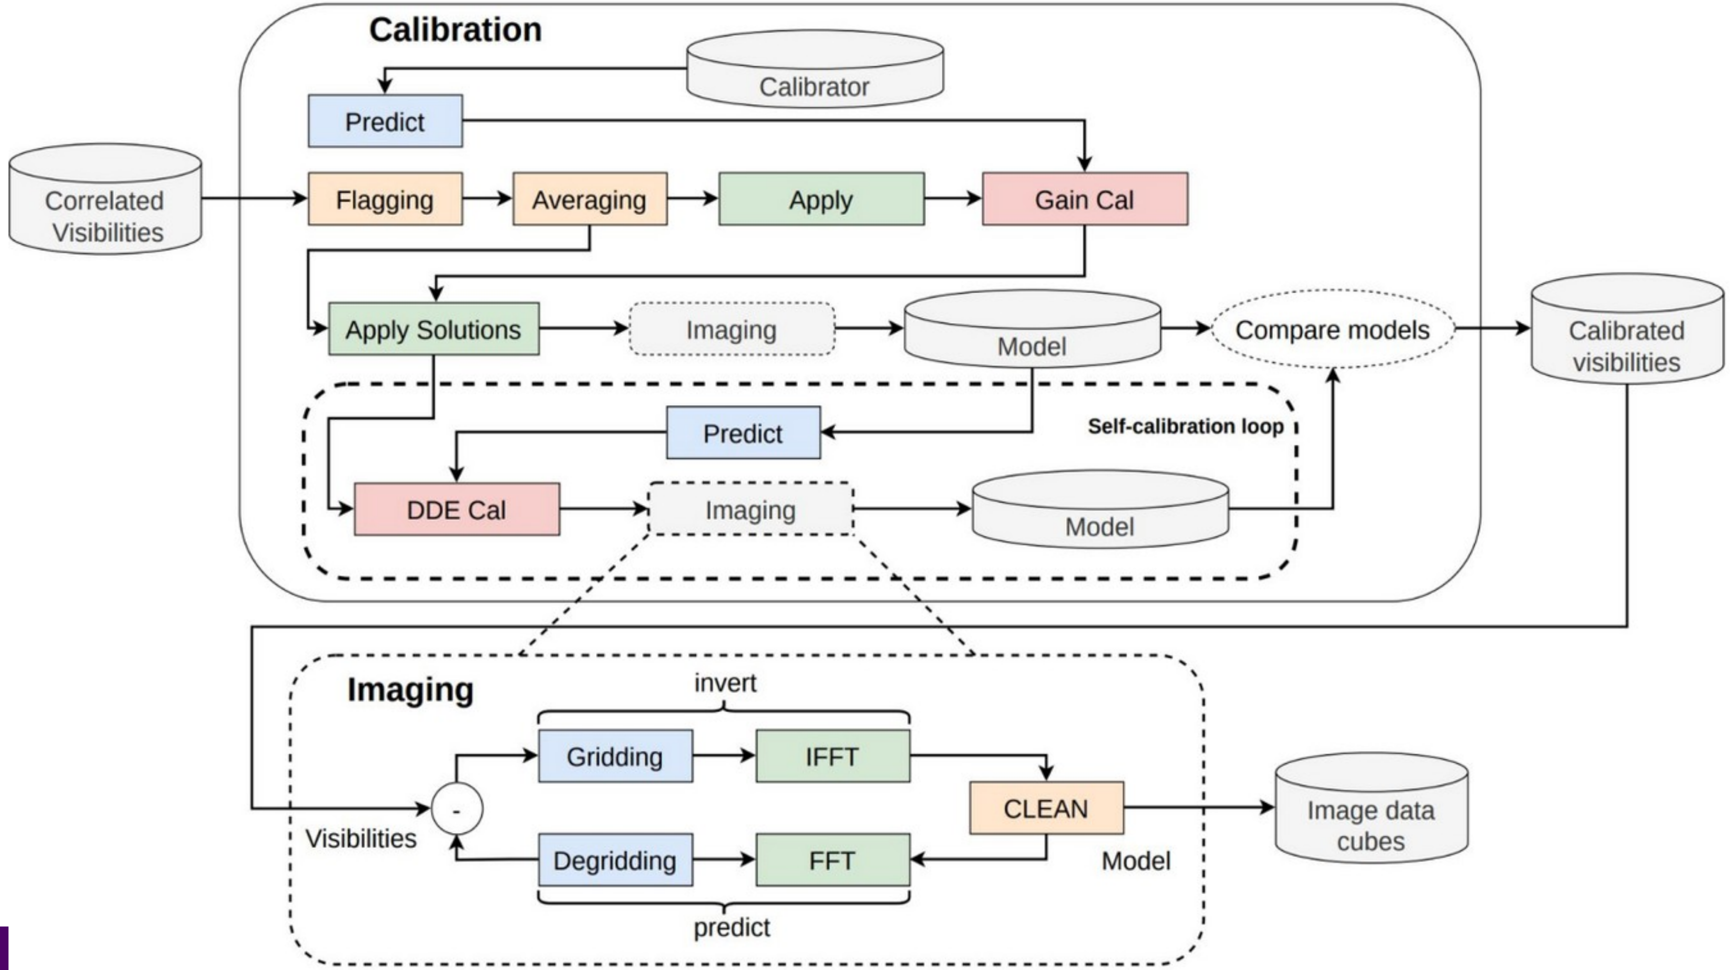

Constraints

Maximum data flow calculation

	Low	Mid
• Raw data production	2 Pbs/s	20 Tb/s
• warrant streaming processing		
• Visibilities		
• consider all pairs	C^{2}_{512} = 130816 baselines	C^{2}_{196} = 19110 baselines
• minimum integration	0.9 s	0.14 s
• maximum flow	7.8 Tb/s	8.9 Tb/s
• require random access but 24h storage is ~ 100 PB/day (including intermediate data)		
• Data products	~100 Gb/s	~100 Gb/s

- Numbers correspond to using full array for a single observation but sub-array and commensal observations modes exist

Total cost of ownership of computing

- Capital
  - procure HW as late as possible to benefit from performance improvements
  - budget is largely underestimated (based on Moore’s law & Dennard scaling, does not account for inflation)
- Operations
  - energy for processing found to be the principal yearly expense
  - tight energy budgets in South Africa and Australia
- Refresh
  - HW refreshed periodically (with potential gains in operations cost)
  - amortized in >5 years

System design

Data processing architecture (I)

- Edge to cloud architecture to manage data flow and complexity

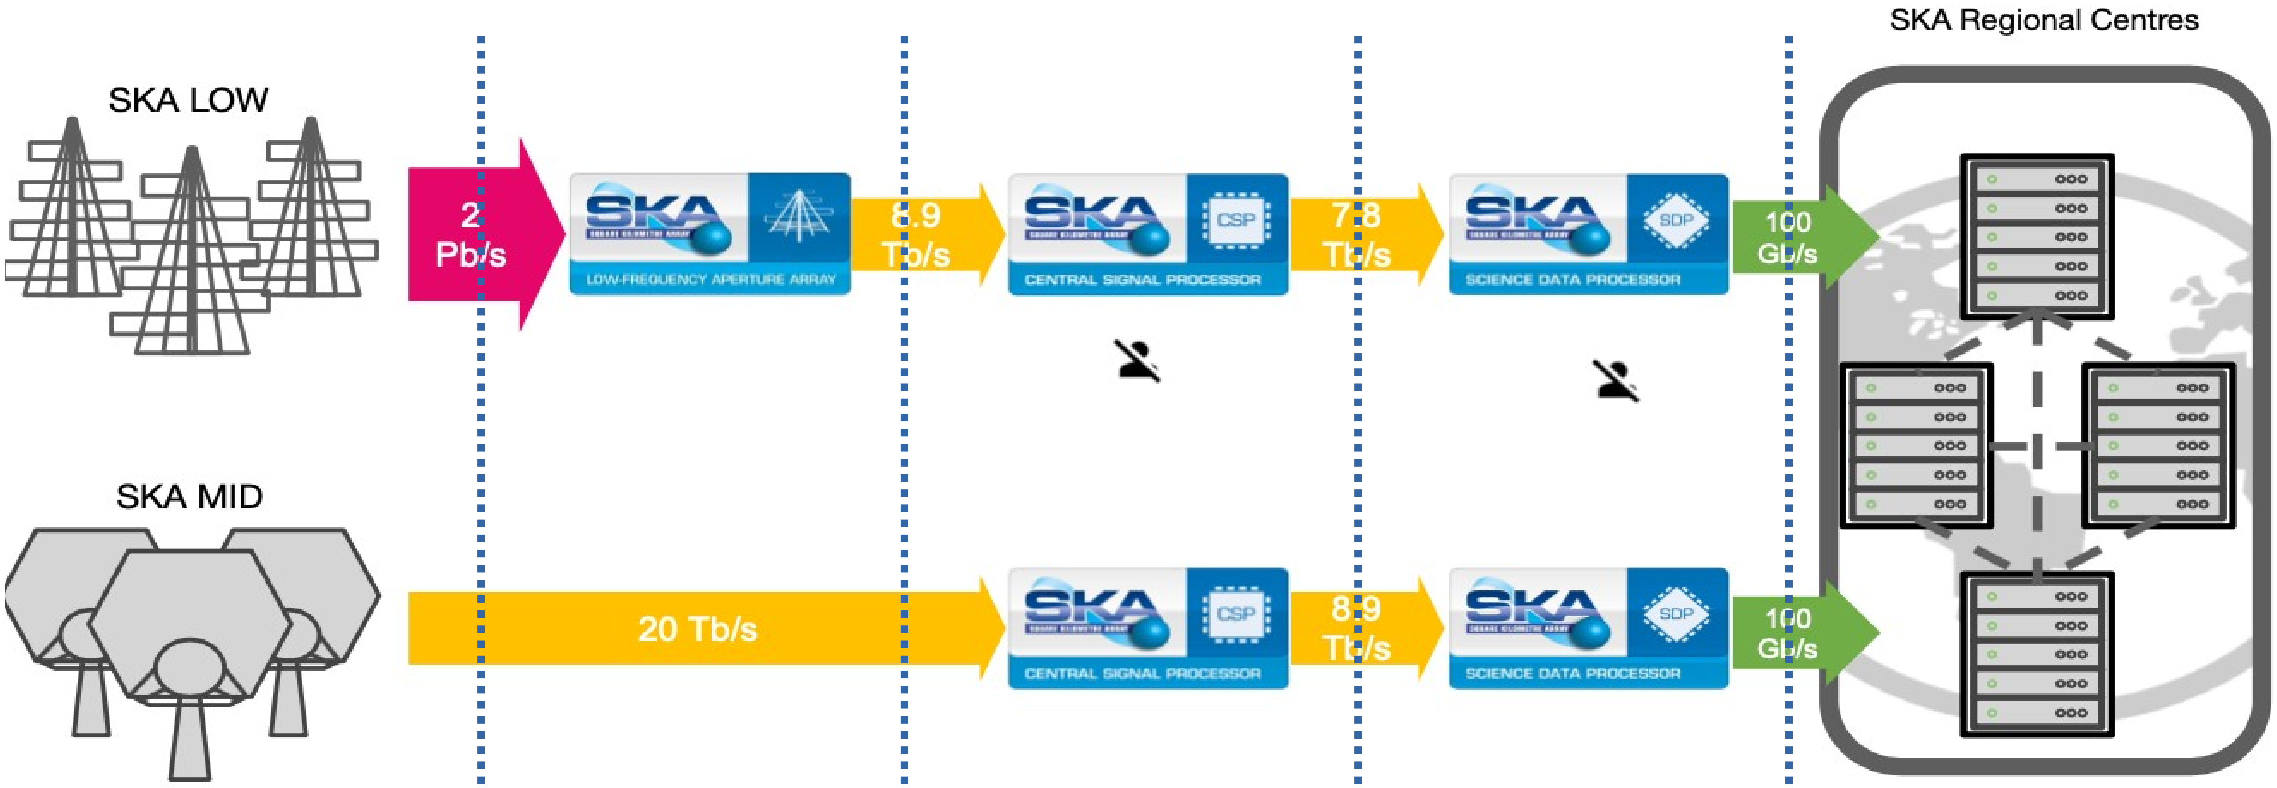

(courtesy SKAO)

Data processing architecture (II)

- Central Signal Processor (CSP)
  - on site in theory, near site in practice (difficulty of installing a Data Processing Center in the desert & radio frequency interferences)
  - specialised system: correlation and beam forming needs are stable, FPGA-based
- Pulsar Search (PSS)
  - off site
  - specialised system: temporal data processing, surveys requires homogeneous products over time, FPGA-based
- Pulsar Timing (PST)
  - off site
  - specialised system: temporal data processing, GPU-based
- Science Data Processor (SDP)
  - off site
  - semi-specialised system
- Science Regional Centers (SRC)
  - international
  - science portal
  - federated cloud-like infrastructure

SDP tasks

- Scientific
  - Alerts
    - detection within a few seconds: high-throughput computing, not demanding computing
    - follow-up: point and observer within a few seconds
  - Pulsar & transient catalogues
  - Image cubes: spatial x frequency channel (x polarisation)
  - Power spectra and line cubes: power versus frequency
  - Raw data (occasionally)

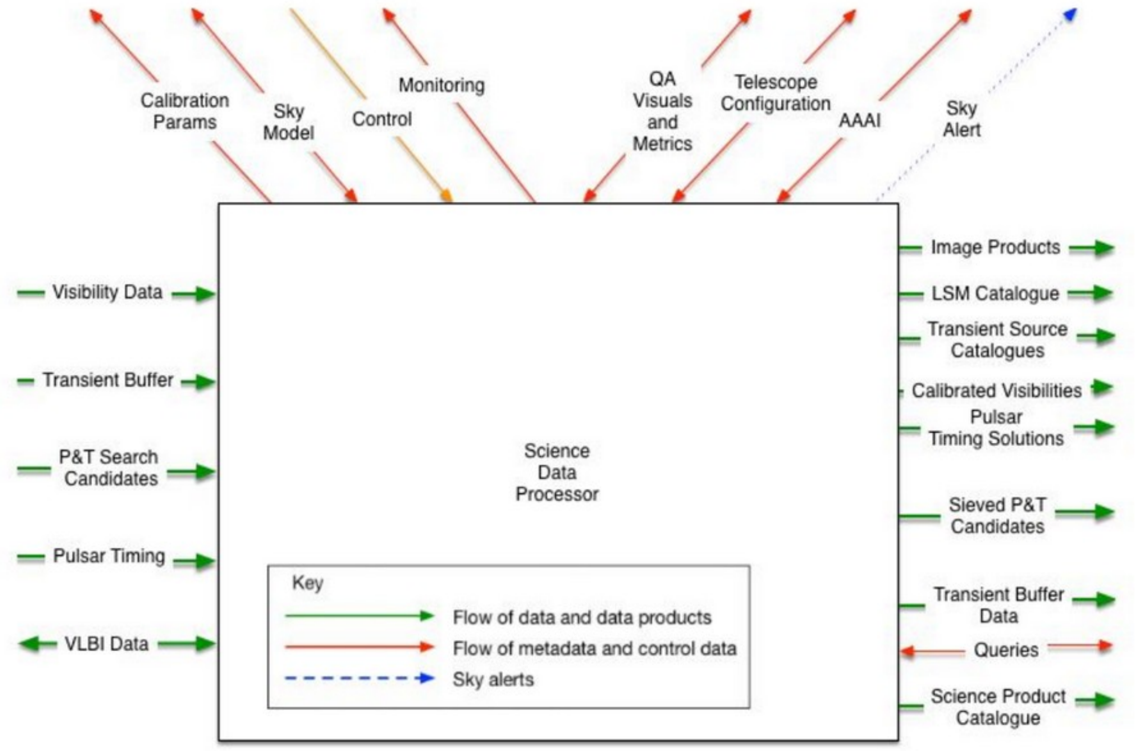

SDP Critical Design Review (courtesy SKAO)

- Internal
  - Calibration: instrument & sky models
  - Closing feedback loops: pointing & beam forming
  - Quality assurance: telescope monitoring & science metrics

Resource estimates

- Parametric model for Low and Mid
  - estimates compute capacity and output data rate
  - based on telescope configuration and average observation schedule
- On average for imaging pipelines only (real-time and batch):
  - Low: 8.47 PFlops (effective), 4.6 Gb/s (without visibilities), 47.4 Gb/s (with visibilities)
  - Mid: 5.29 Pflops (effective), 28.4 Gb/s (without visibilities)
- Does not account for other tasks on SDP
- Theoretical model which does not account for variable complexity (sky, interference, content of field of view)
- Peak capability estimates ~120 PFlops (~10^{th} rank in Top500) with high efficiency estimate (for peak to effective conversion)

AA2 scaling challenge

- AA2 matches the largest radio-telescopes to-date (LOFAR, MeerKAT)
- Objectives by 2026
  - begin producing scientific results
  - validate design by processing AA2 data at its acquisition rate
  - AA* is 2 years later with numbers of visibilities increased by ~5 (Mid) and ~23 (Low)
- Track ICAL pipeline development
  - existing SW does not scale
  - Current status of benchmarks (extrapolated from smaller runs): Low x7.47, Mid: x747 from goal

Estimated SDP Scaling: AA1→AA4

(~50x in 17 months! Qualitative only, underestimates the AA2 situation)

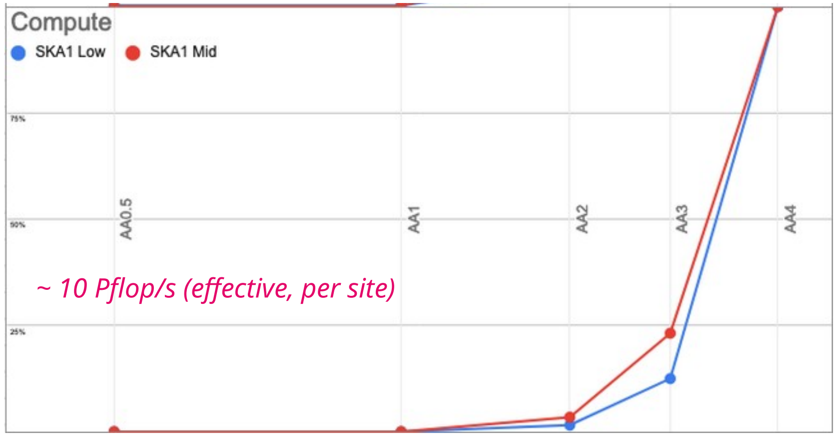

Compute ramp up (courtesy Peter Wortmann)

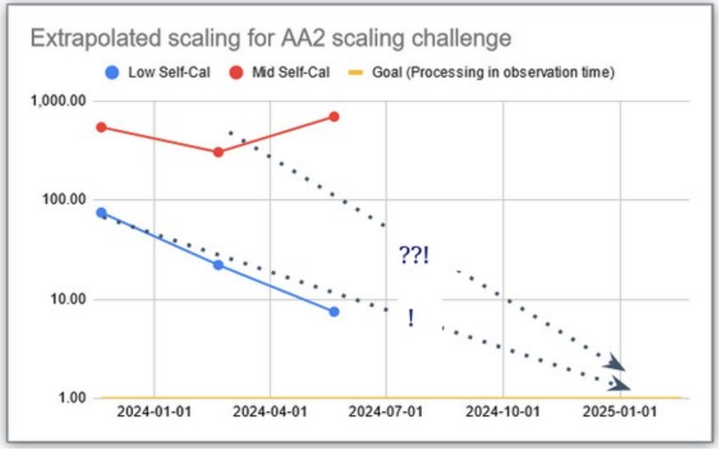

Performance improvement (courtesy Peter Wortmann)

Power

- Science Processing Centers in Cape Town and Perth
  - most of CSP (after AA2)
  - SDP
  - Infrastructure & Cooling
- Cost & greenhouse gas issue: multiple power sources (grid/solar/diesel/battery/RUPS)
- SDP allocation
  - average: 1.3 MW Mid / 1.6 MW Low
  - peak: 2.0 MW Mid / 2.23 MW Low
- Green500: Frontier, Lumy, Adastra achieve ~100 PFlops @ 2MW at maximum efficiency

SKA1-Mid SPC/ SOC Power Budget in Cape Town				
Products	*AA4 Long Term Average (>30min) [kW]*	*AA4 Peak Instantaneous (<5sec) [kW]*	*AA* Long Term Average (>30min) [kW]*	*AA* Peak Instantaneous (<5sec) [kW]*
PDT4 - MID Digitisation	**230.8**	**323.4**	**230.8**	**323.4**
CSP.CBF	230.8	323.4	230.8	323.4
PDT6 - Network & Computing	**1641.7**	**2481.9**	**589.3**	**872.6**
SDP Hardware MID	1300.0	2000.0	325.0	500.0
PSS Hardware MID	296.0	414.0	222.0	310.5
PST Hardware MID	16.4	26.8	12.3	20.1
OMC Hardware MID	12.9	18.1	12.9	18.1
NSDN MID	5.6	7.8	6.6	9.2
CPF-SPC link MID	8.2	11.5	7.9	11.1
NMGR	2.6	3.6	2.6	3.6
Building losses and cooling	**374.5**	**561.1**	**164.0**	**239.2**
Commissioning Margin	**224.7**	**336.6**	**98.4**	**143.5**
Site Total	**2471.7**	**3702.9**	**1082.6**	**1578.8**

SKA1-Low SPC/ SOC Power Budget in Perth				
Products	*AA4 Long Term Average (>30min) [kW]*	*Peak Instantaneous (<5sec) [kW]*	*AA* Long Term Average (>30min) [kW]*	*AA* Peak Instantaneous (<5sec) [kW]*
PDT6 - Network & Computing	**1629.2**	**2270.9**	**429.2**	**598.4**
OMC Hardware LOW	12.9	18.1	12.9	18.1
SDP Hardware LOW	1600.0	2230.0	400.0	557.5
NSDN LOW	6.5	9.1	6.5	9.1
CSP-SDP LOW	7.2	10.1	7.2	10.1
NMGR	2.6	3.6	2.6	3.6
Building losses and cooling	**325.8**	**454.2**	**85.8**	**119.7**
Commissioning Margin	**195.5**	**272.5**	**51.5**	**71.8**
Site Total	**2150.6**	**2997.6**	**566.6**	**789.9**

Fall-back options

- Risk: if processing does not keep up, the storage is filled and visibility data cannot be ingested any more
- Scenarios
  - cancel up-coming observation: do not observer continuously (as all existing radio-telescopes have at first light)
  - cancel on-going jobs and discard associated data (free resources for the next observation)
  - push data towards the SRCNet
    - existing radio-telescopes transfer data to users
    - more complex for the SKA due to the data flow: intercontinental network capacity prohibits transferring raw visibilities
- Frugality reinvented: produce little at an extreme cost

The value of visibilities

Managing visibilities

- Visibilities drive complexity
  - manage resources via adaptively handling visibilities
  - already done routinely as not all tasks require maximum resolution
- Time and baseline averaging are part of pre-processing
  - preserve gain from integration
  - induce smearing as frequential information is degraded ((u, v) coverage)

Processing visibilities

- Weighting schemes are routinely used
  - control the importance of baselines
    - short baselines are numerous: high SNR, low angular resolution
    - long baselines are few: noisy data, high angular resolution
  - impact on SNR and PSF (conditions deconvolution)
  - still all visibilities are processed
- Key is to adaptively process the visibilities
  - on-going developments at OCA for partitioning the baselines
    - simplifies the reconstruction problem for each baseline class
    - produce intermediate results for compressing visibilities

On-going efforts

- SKA France in kind contribution to SKA for sustainable computing
- SKA France & DDN involved in co-design studies with a current focus on procuring AA2 HW
- SKA chosen as an illustrator for Exa-DoST (NumPEx PC3)
  - evaluate I/O and storage management libraries
  - efficiency requires more than just optimising storage (compute, scheduling etc.) find the right balance: coordinate with Exa-Ma, Exa-AToW and energy
- Phase 0 for constructing the French SRC node: report to the MESR in October 2024 for funding
  - at least 3% of SRCNet resources to comply with SKAO rules
  - support the French community's use of SKA data and precursors (LOFAR, NenuFAR)
  - 4 working groups: governance model for a distributed infrastructure, HW, SW and science roadmap

We recognise and acknowledge the Indigenous peoples and cultures that have traditionally lived on the lands on which our facilities are located.

www.skao.int