SKA and its computing challenges

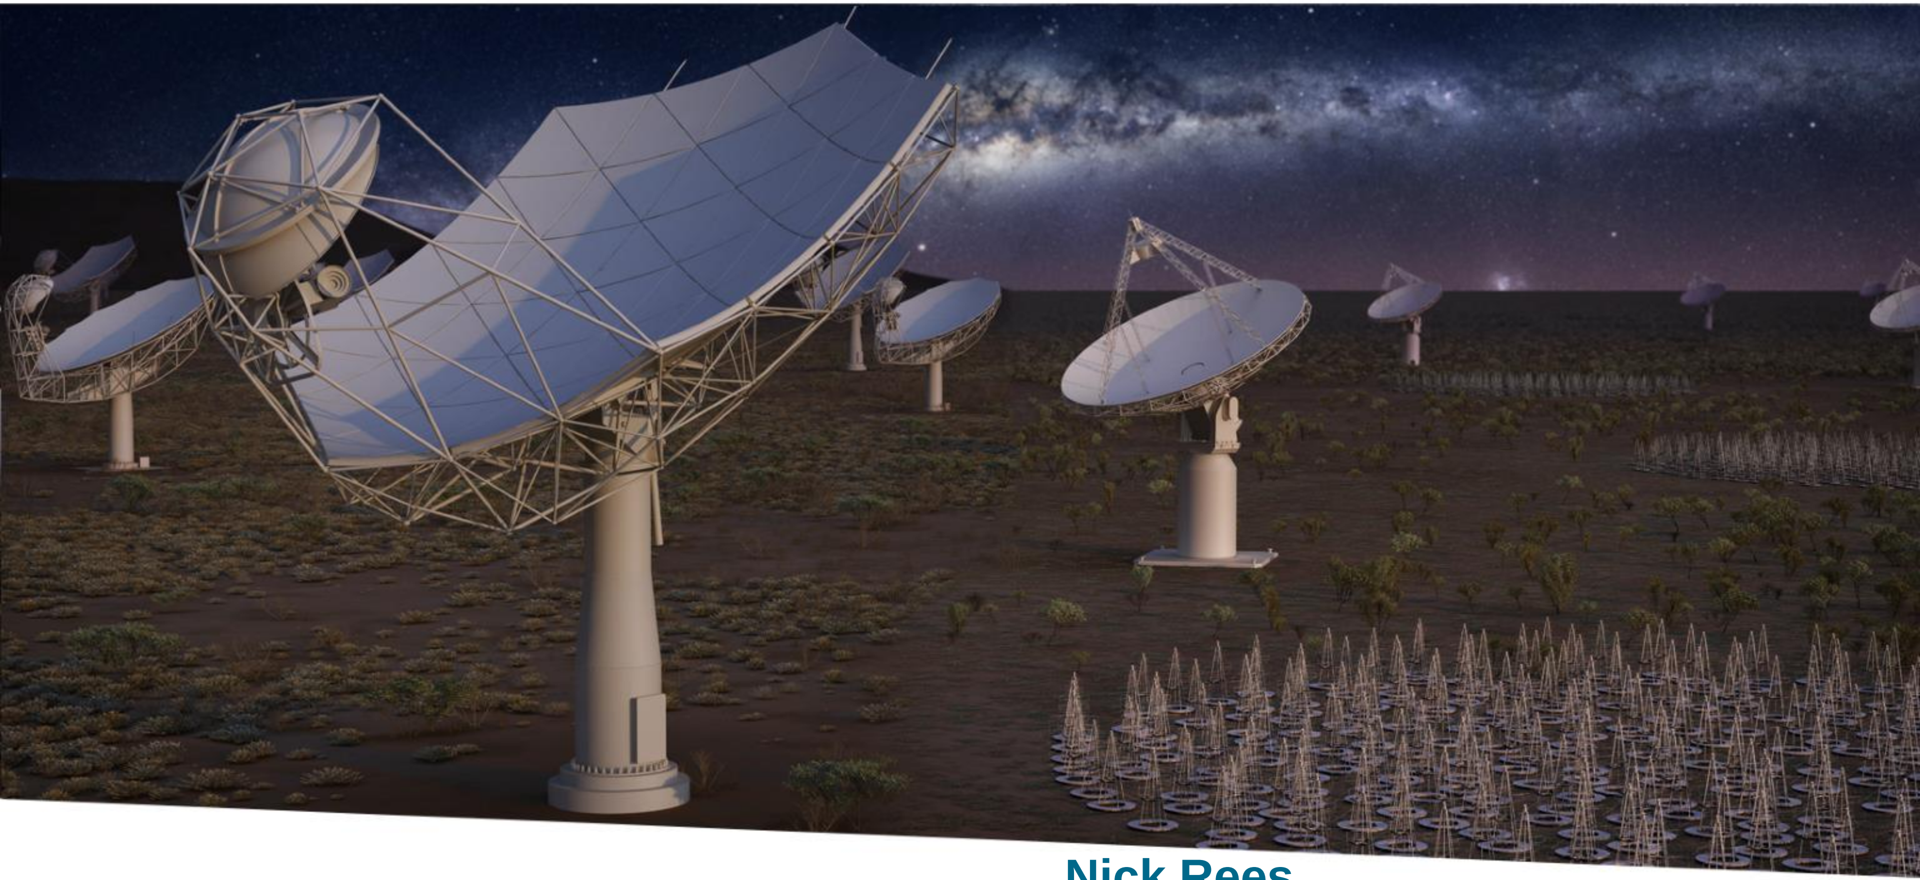

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope

Nick Rees

Head of Computing and Software

16th May 2017

Summary

- The Project
- The Science
- Computing and Software
  - Telescope Manager
  - Low Frequency Aperture Array
  - Central Signal Processor
  - Science Data Processor
- Regional Centres
- Conclusions

The Project

Square Kilometre Array

3 sites; 2 telescopes + HQ

1 Observatory

Design Phase: > €170M; 600 scientists+engineers

Phase 1

Construction: 2018 – 2024

Construction cost cap: €674.1M (inflation-adjusted)

Operations cost: under development (see below)

MeerKat integrated

Observatory Development Programme (€20M/year planned)

SKA Regional centres out of scope of centrally-funded SKAO.

Phase 2: start mid-2020s

~2000 dishes across 3500km of Southern Africa

Major expansion of SKA1-Low across Western Australia

HQ in UK; telescopes in AU & SA

- SKA1-LOW: 50 – 350 MHz
- Phase 1: ~130,000 antennas
- across 65km

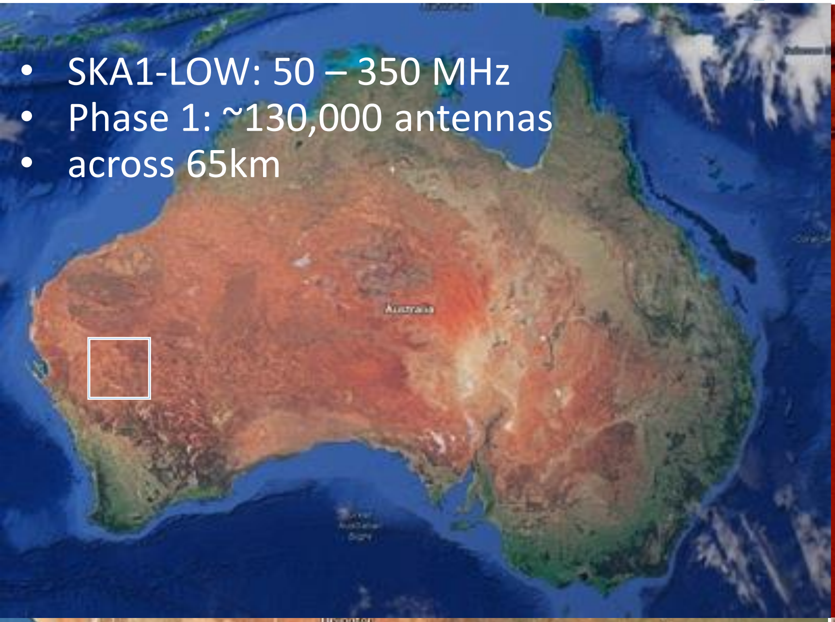

SKA1-Mid: 350 MHz – 24 GHz
Phase 1: 200 15-m dishes across 150 km

Construction: 2018 – 2024; Cost cap: €675M

SKA Design Consortia

SIGNAL AND DATA TRANSPORT

TELESCOPE MANAGER

ASSEMBLY, INTEGRATION & VERIFICATION

DISH

SQUARE KILOMETRE ARRAY

SCIENCE DATA PROCESSOR

LOW-FREQUENCY APERTURE ARRAY

CENTRAL SIGNAL PROCESSOR

INFRASTRUCTURE AUSTRALIA

INFRASTRUCTURE SOUTH AFRICA

Precursors

SKA Organisation: 10 countries, more to join

Australia (DoI&S)
Canada (NRC-HIA)
China (MOST)
India (DAE)
Italy (INAF)
Netherlands (NWO)
New Zealand (MED)
South Africa (DST)
Sweden (Chalmers)
UK (STFC)

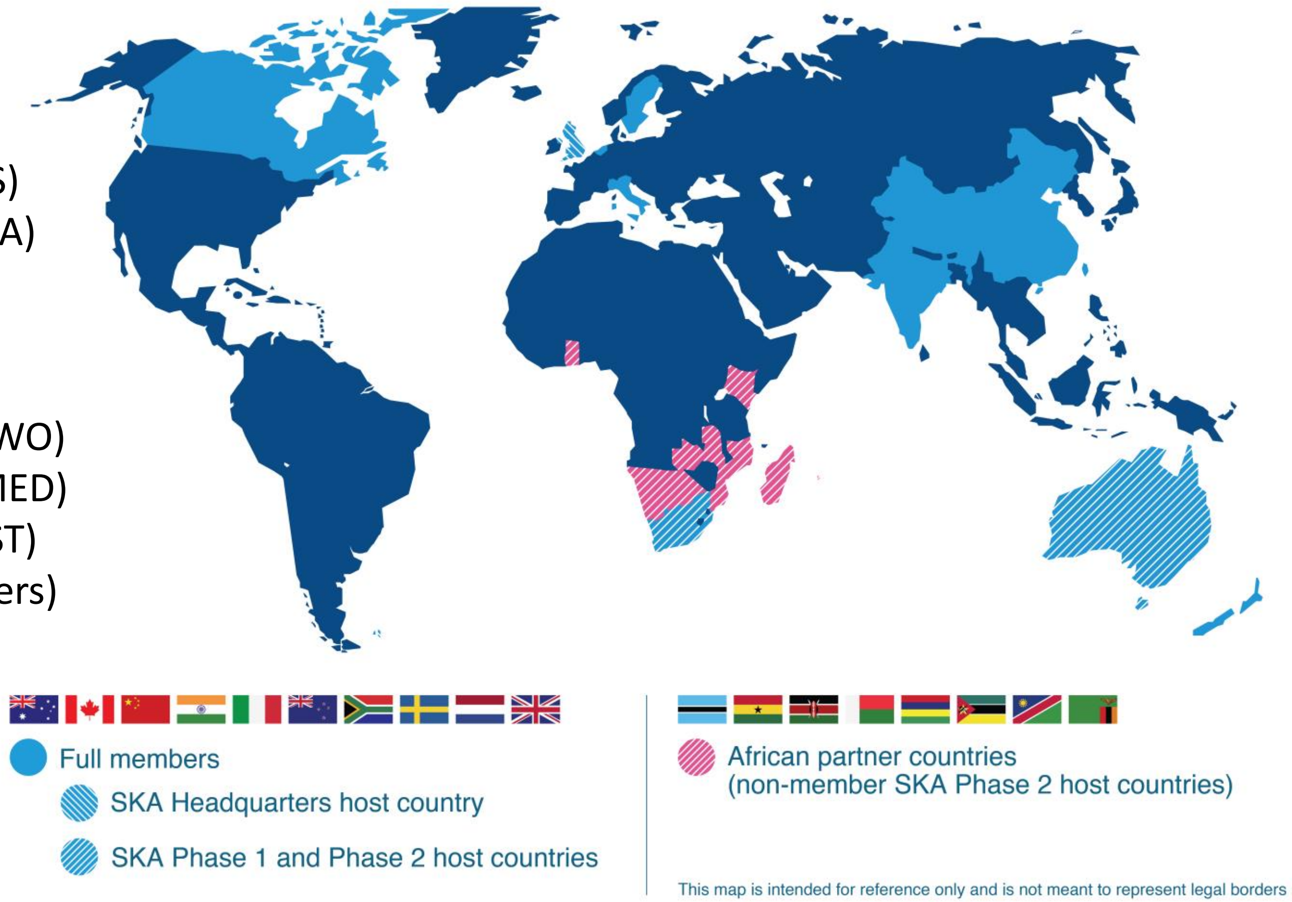

Interested Countries:

- France
- Germany
- Japan
- Korea
- Malta
- Portugal
- Spain
- Switzerland
- USA

Contacts:

- Mexico
- Brazil
- Ireland
- Russia

Status of interested countries

- Portugal: Letters of support from Ministers of Science and Economy. Announcement imminent.
- Germany:
  - MPG providing funding for a second SKA1-Mid prototype dish (first to go to site)
  - Germany attended IGO meetings following positive re-engagement.
- France: Accelerated re-examination of Astrophysics section of National Science Infrastructure Roadmap. Engagement with industrial partners.
- Spain: Spanish State Secretary has written to D-G, supportive of joining SKA in near future .
- Switzerland: Swiss State Secretary requested observer status at SKA Board - granted; indication will join when an IGO.
- Japan: Attended SKA Board in November and March.
- Korea: Attended SKA Board in March.
- USA: Establishing radio astronomy strategy for Astro2020. Ongoing discussions with Director NRAO and DoE labs.

Future SKA governance structure

- IGO = ‘Convention’ agreed between governments
  - Government commitment: Long-term political stability, funding stability
  - A level of independence in structure
  - Availability of ‘supporting processes’ through Privileges and Immunities from members: functional support for project
  - ‘Freedom to operate’, specifically through procurement process, employment rules and so on

Operations Scope

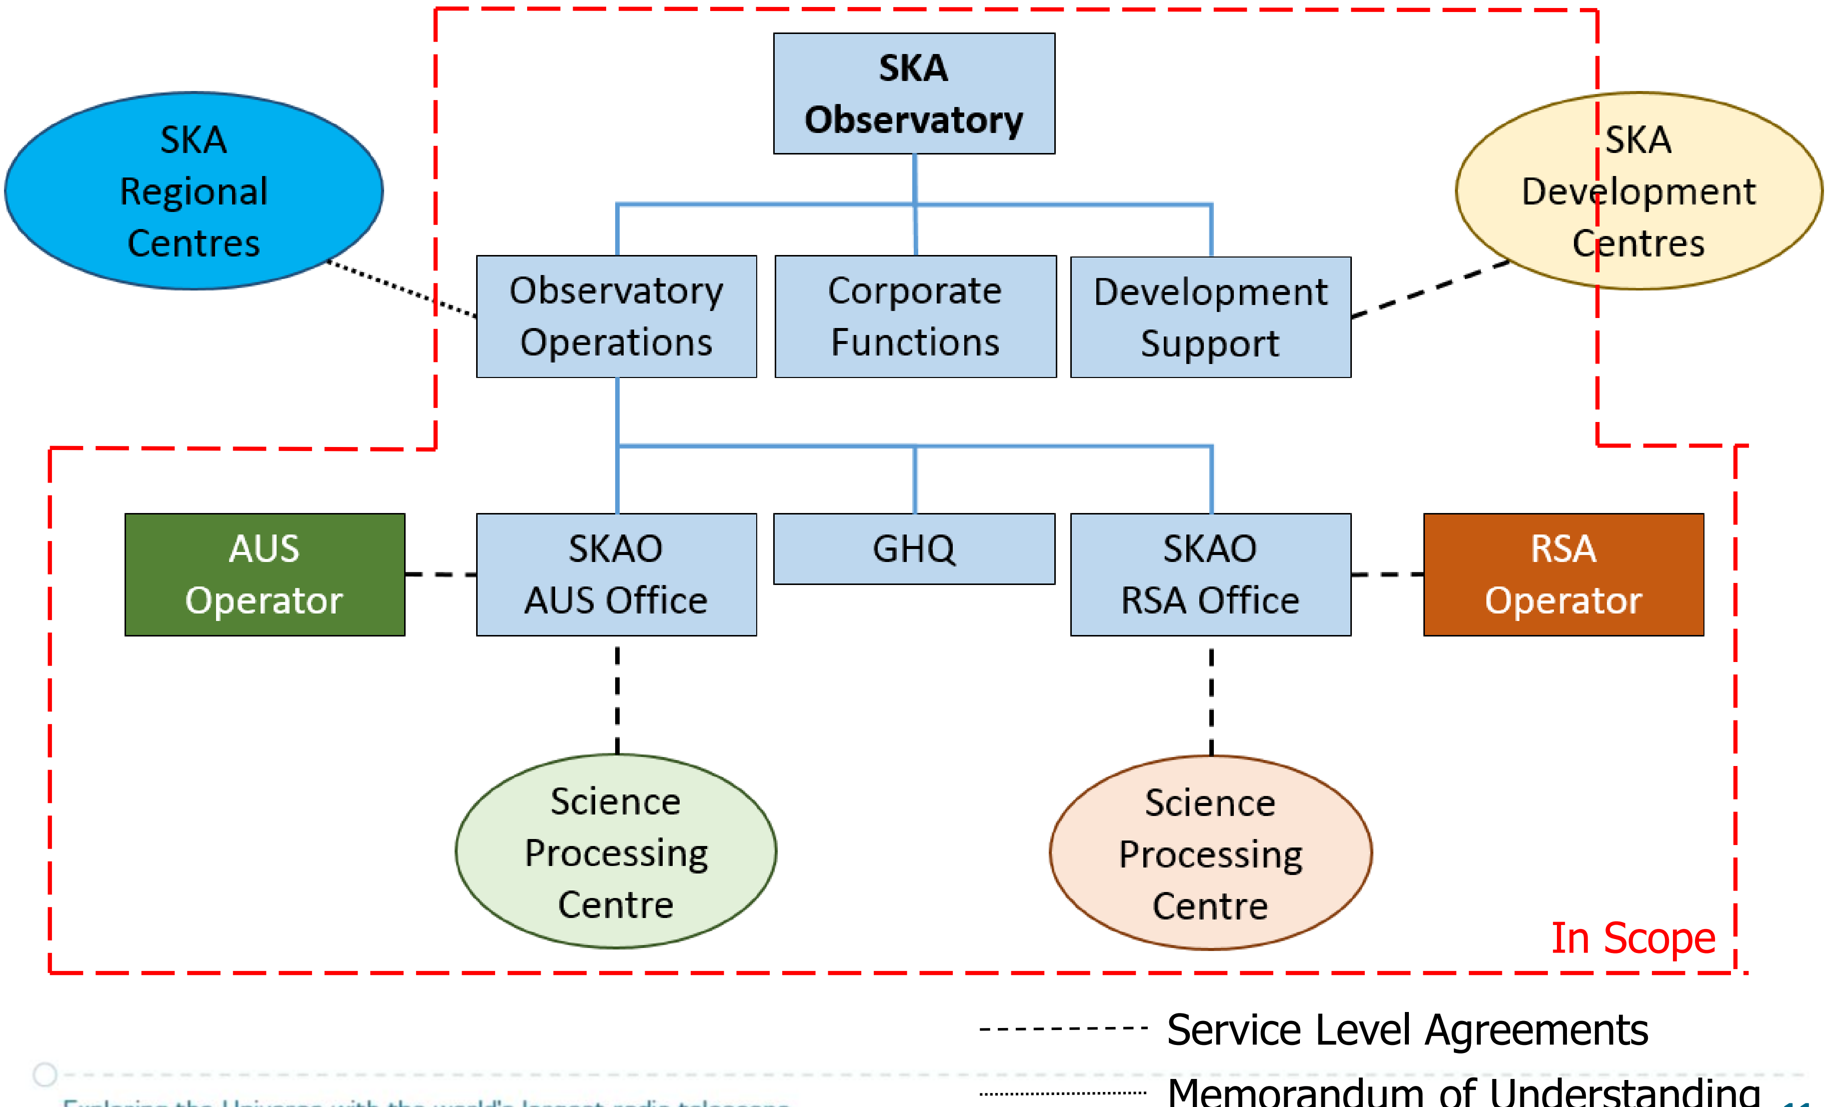

Overall project timeline – to be confirmed

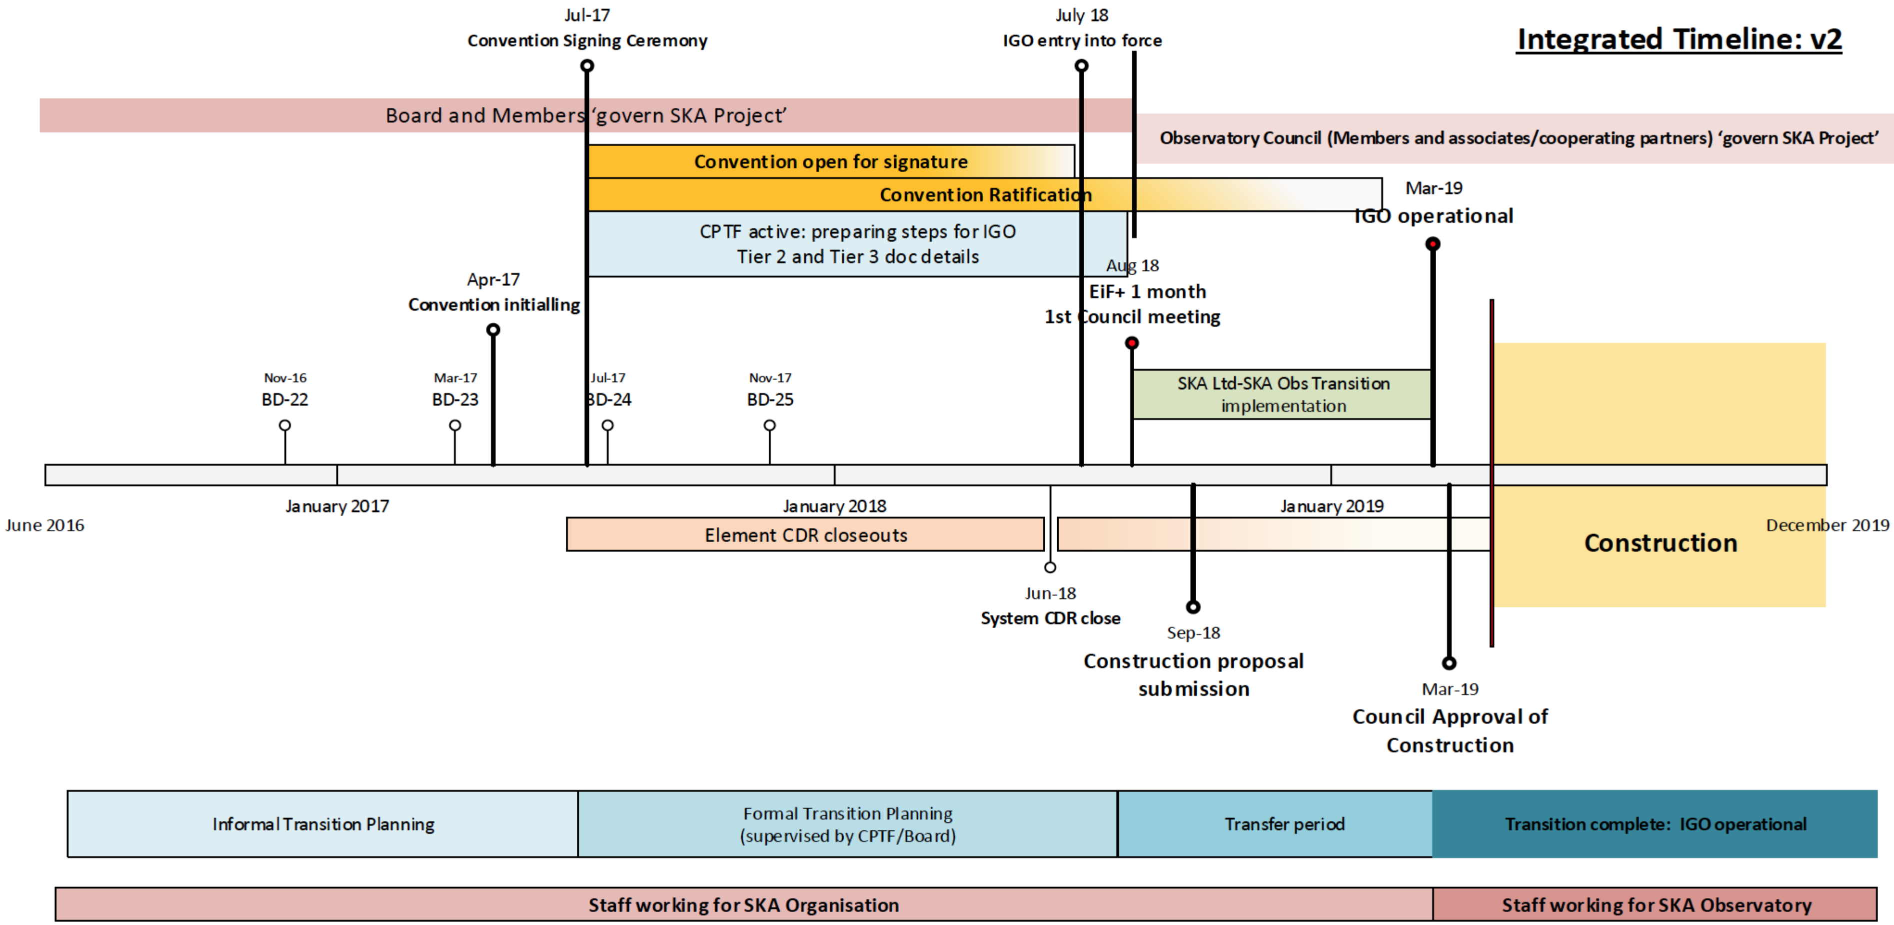

Key dates:

- Convention signing July 2017
- CDRs Q4 2017 – Q2 2018
- IGO enter into force July 2018
- SKA1 Construction approval early 2019

The Science

SKA– Key Science Drivers:
The history of the Universe

Extremely broad range of science!

Era of Recombination

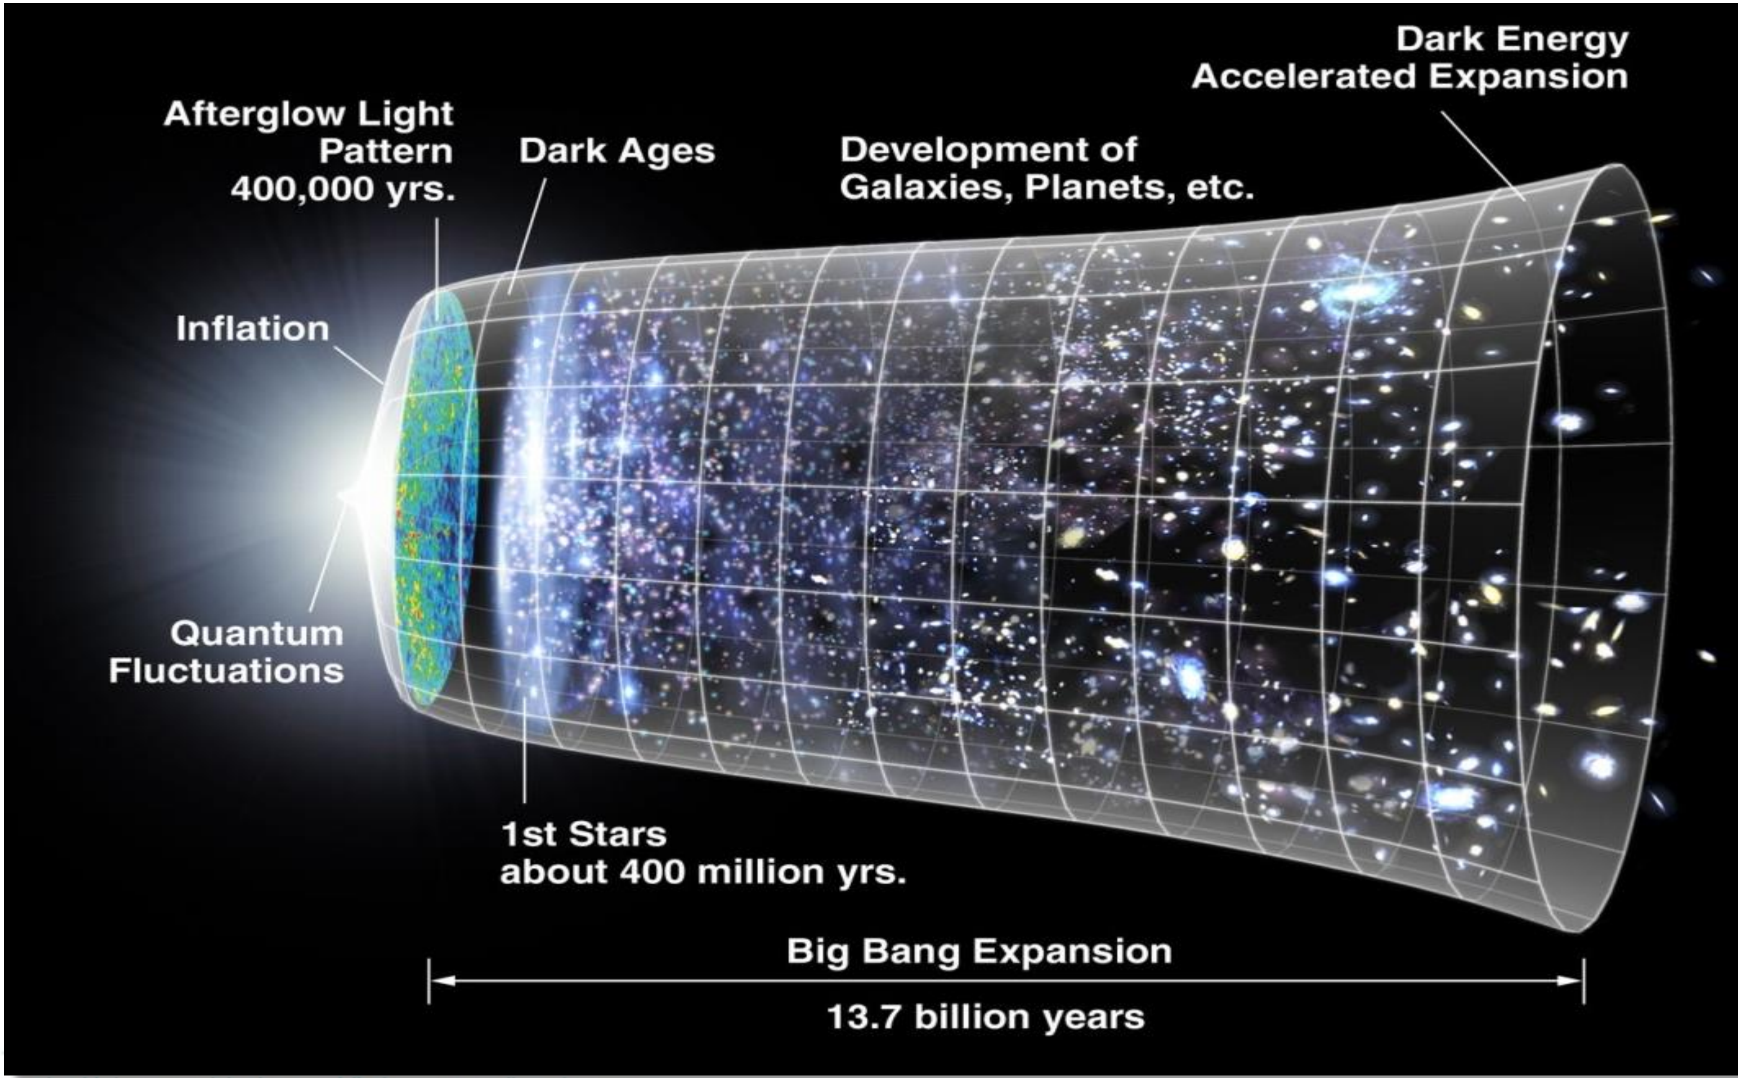


Cosmic Microwave Background
A
Cosmic Dark Ages
B
First Stars
First Galaxies & Black Holes
C
D
Hot Bubble Dominated Epoch
E
Modern Galaxies Form

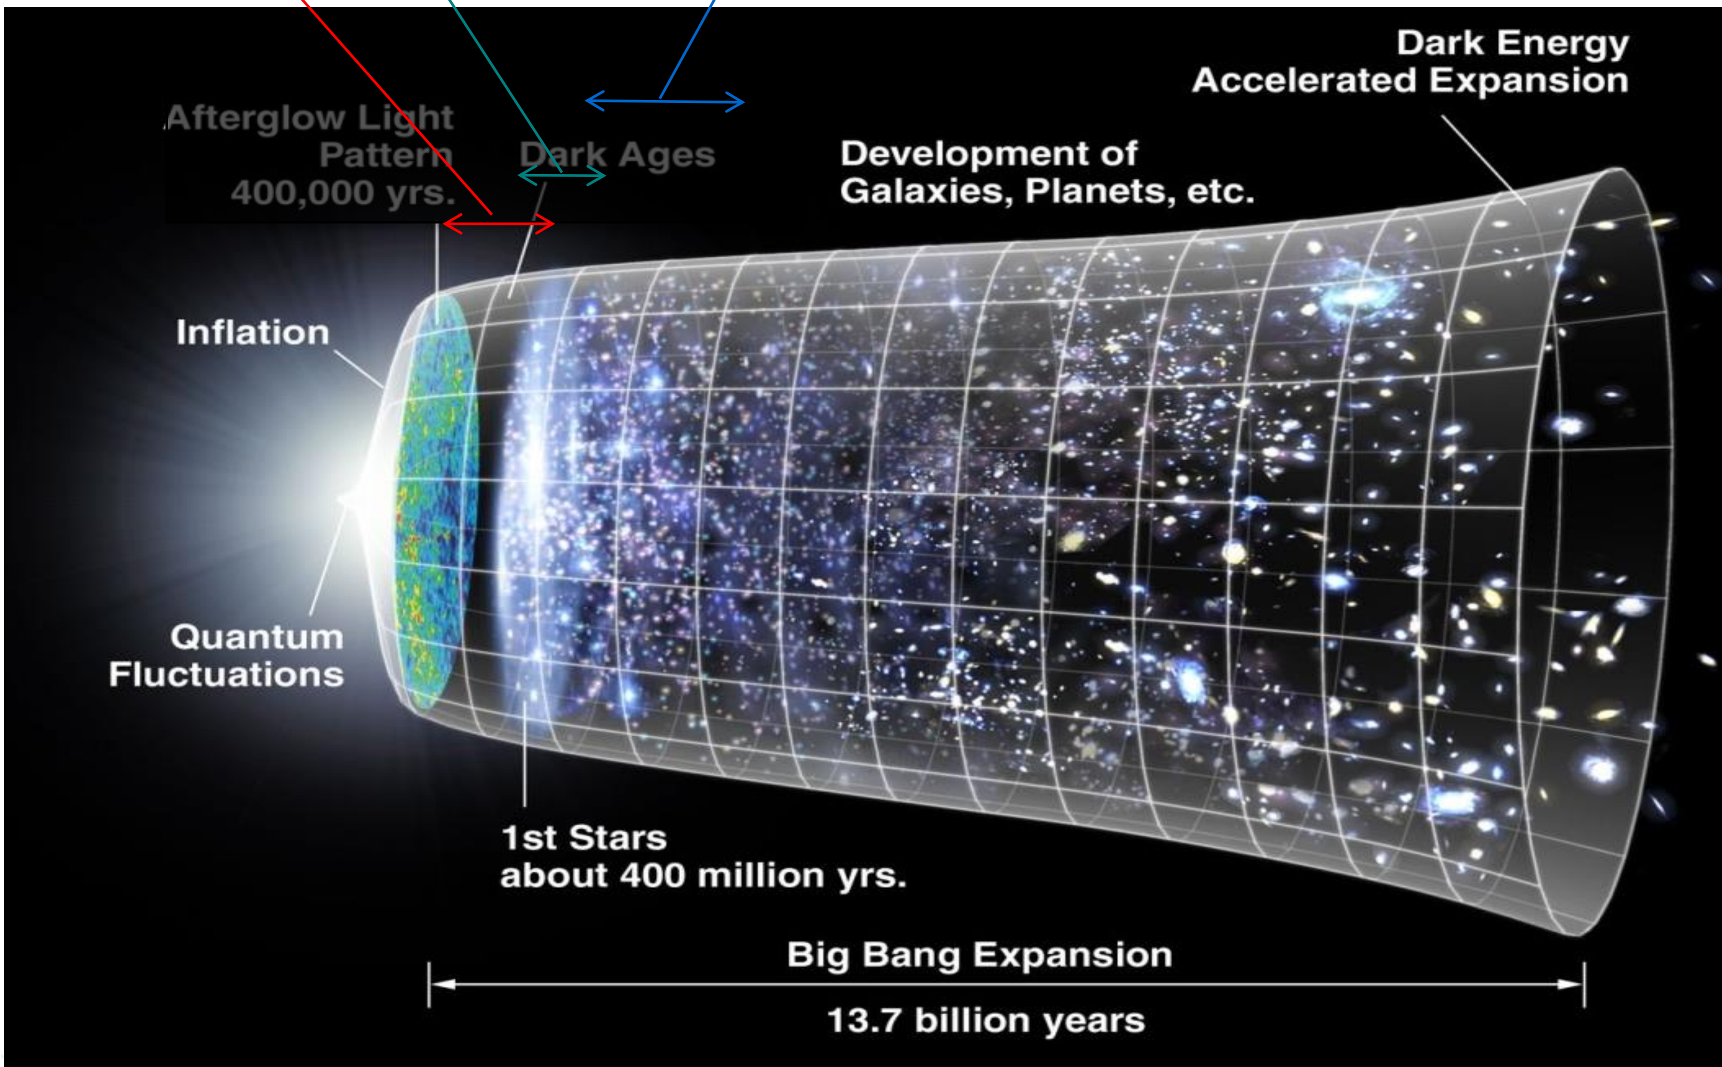
Dark Energy Accelerated Expansion
Afterglow Light Pattern 400,000 yrs.
Dark Ages
Development of Galaxies, Planets, etc.
Inflation
Quantum Fluctuations
1st Stars about 400 million yrs.
Big Bang Expansion
13.7 billion years

SKA1 capability vs state-of-the-art

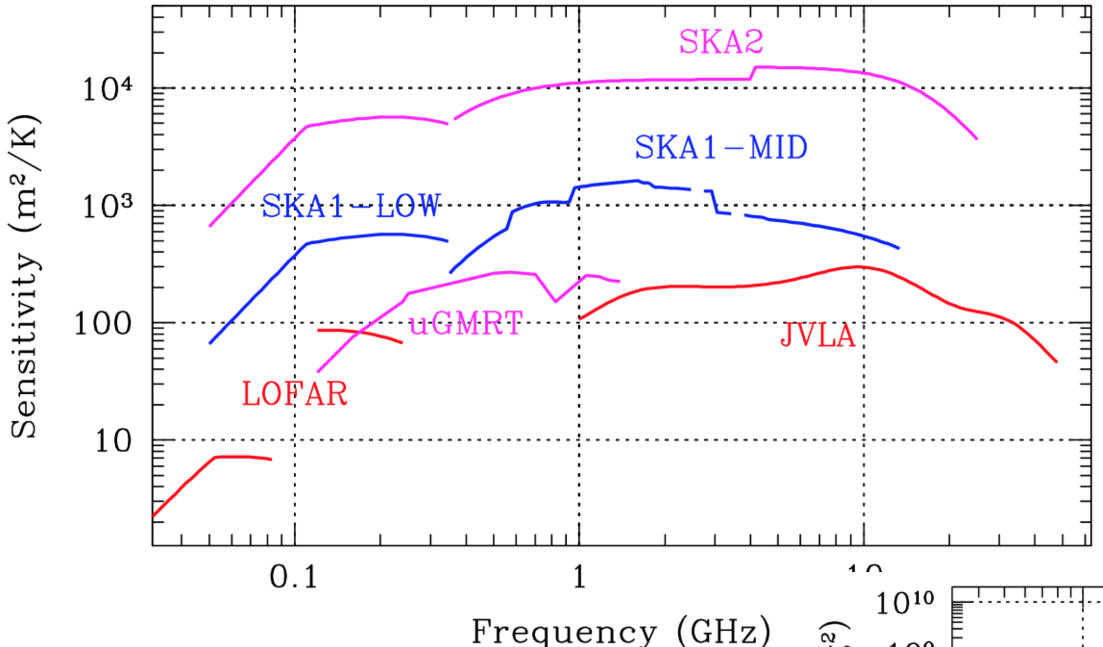

Point-source sensitivity:
~ 4 – 20 times state-of-the-art

Survey speed:
~ 10 - 100 times state-of-the-art

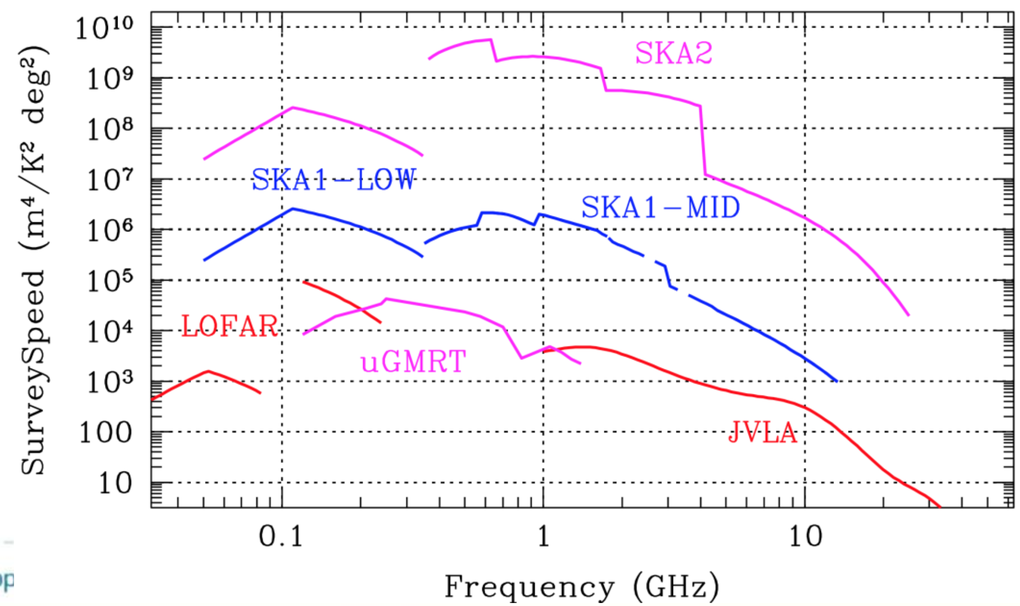

Image Quality Comparison

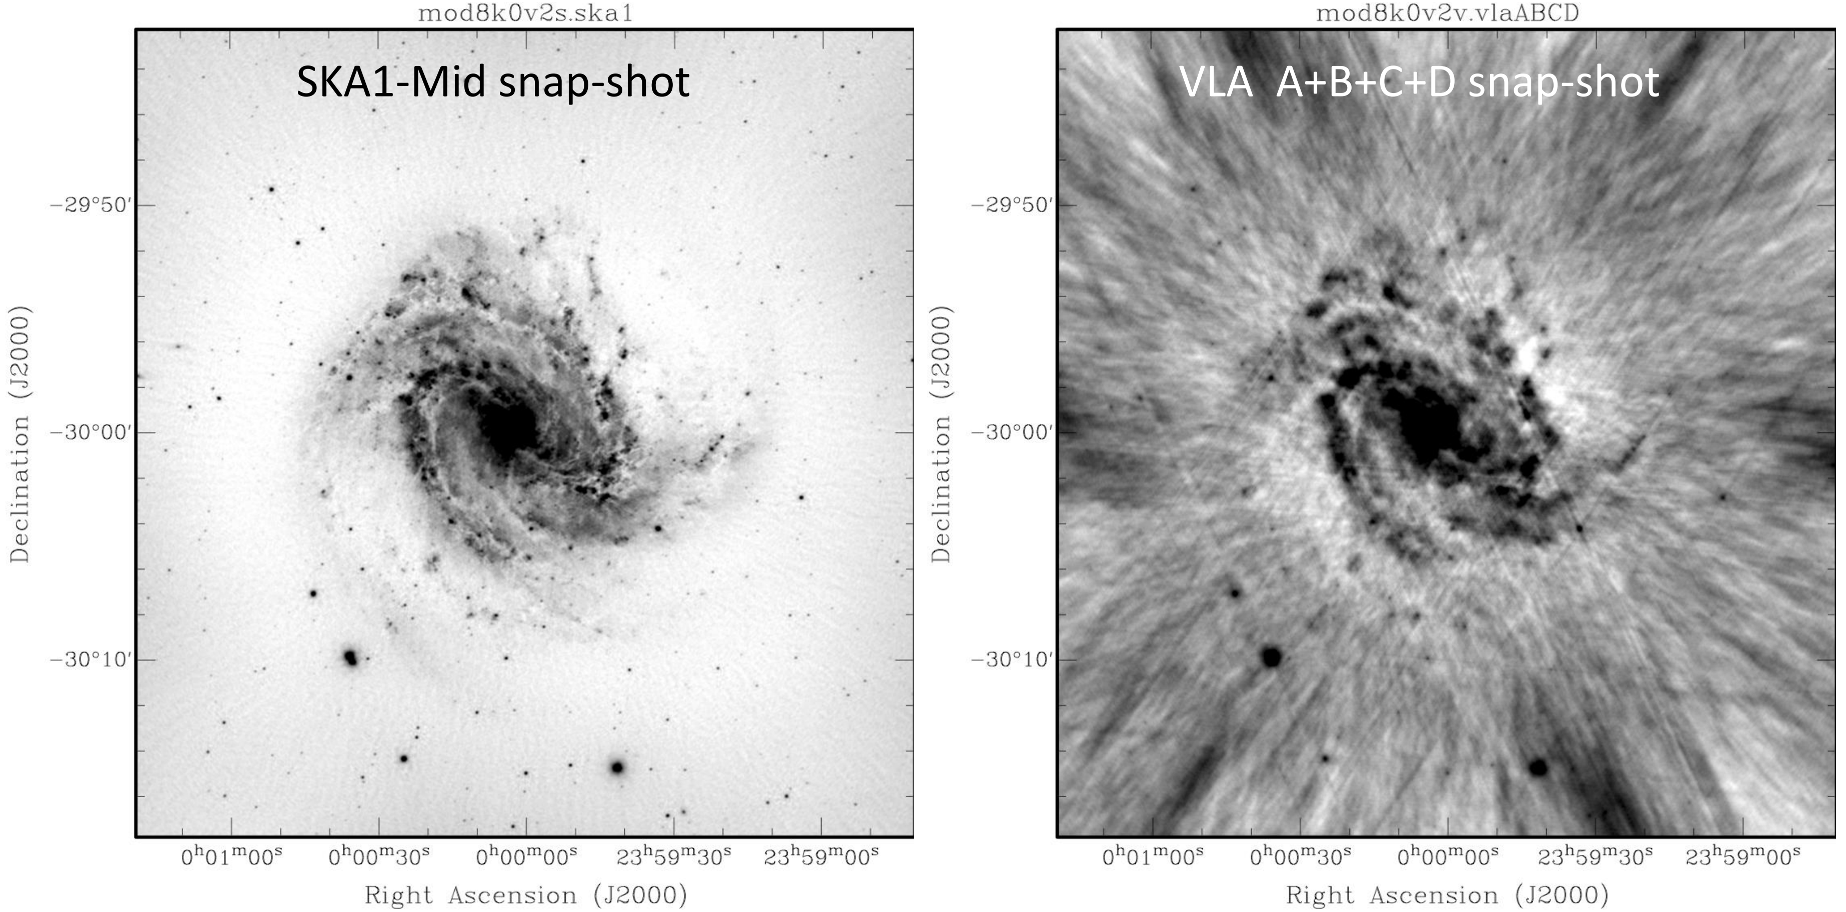

- Single SKA1-Mid snap-shot compared to combination of snap-shots in each of VLA A+B+C+D

Science Working Groups

- Primary scientific community interface to the SKA
- Current SWGs represent a wide range of scientific areas:
    - Extragalactic Spectral Line (non-HI)
    - Our Galaxy
    - Solar, Heliospheric & Ionospheric Physics
    - Epoch of Reionization
    - Cosmology
    - Extragalactic Continuum (galaxies/AGN, galaxy clusters)
    - Cradle of Life
    - HI galaxy science
    - Magnetism
    - Pulsars
    - Transients
- Technique focused Working Group:
    - VLBI
- Topical Focus Group:
    - High Energy Cosmic Particles

Key Science Projects

- Key Science Projects (KSPs) are the science community's highest priority science objectives which are:
  - Consistent with capabilities of the SKA1 design
  - Consistent with a realistic observing schedule filled at 50 – 75% for the first 5 years of scientific operations
- KSP policy currently progressing in context of IGO negotiations
  - Total access (sum of KSP + PI projects) approximately proportional to country's contribution.
  - Mix of KSP/PI projects up to individual member countries
- Call for KSP proposals will happen during construction
  - Will need to allow time for organization and resourcing
  - Is anticipated to lead to large, multi-national teams.

SKA Computing and Software

Introduction

- SKA is a software telescope
  - Very flexible and potentially easy to reconfigure
  - Major software and computing challenge
- Computing challenges are significant
  - Science Data Processor (SDP) needs 25 PetaFLOPS/sec of delivered processing
    - Current estimate is that SDP needs 250 PFLOP/sec peak.
    - Tianhe-2 – 50 PetaFLOPS/sec peak.
    - Memory bandwidth is ~200 PetaBytes/sec
  - Pulsar Search is an additional 50 PFLOP/s of peak processing
  - Power efficiency required is ~40x better than Tianhe-2,
- Software challenges are also large
  - Feb 2017 costings have ~€90M of software.

System Overview

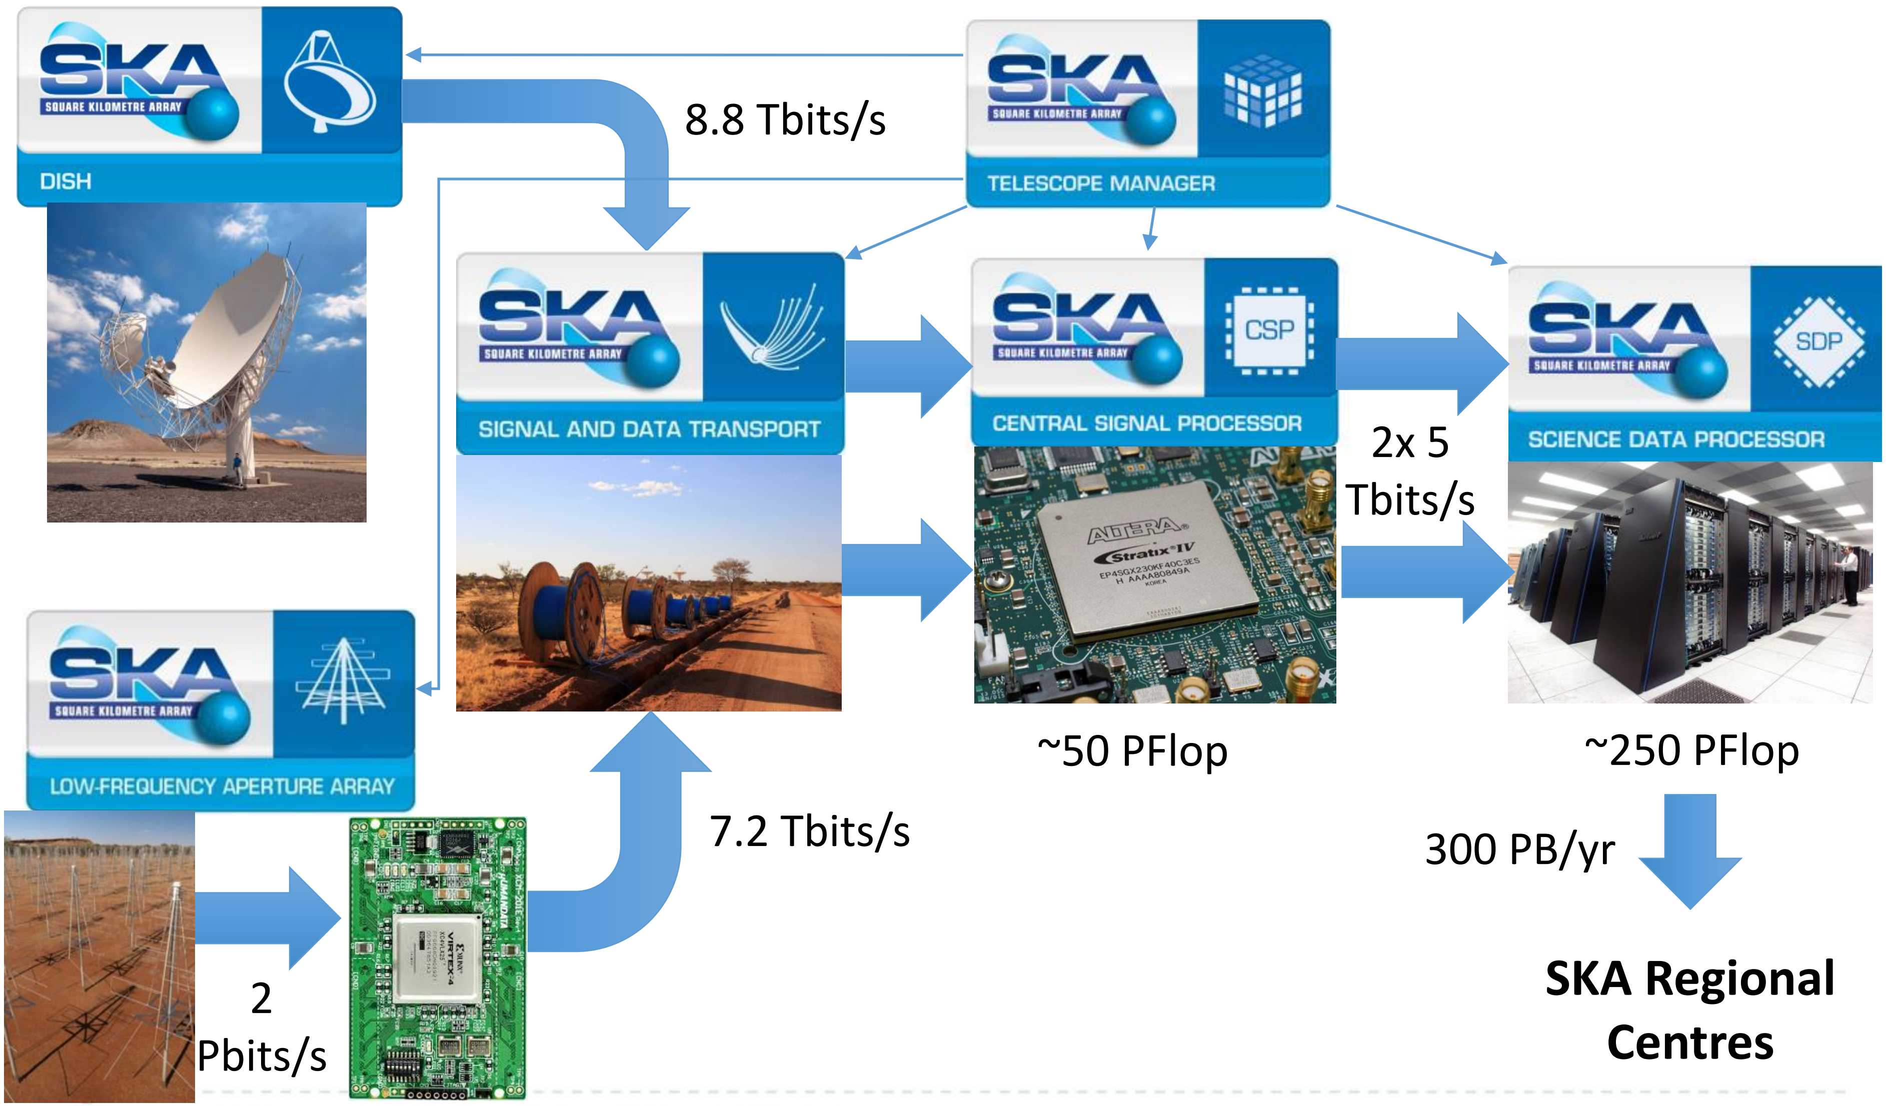

Telescope Manager

Telescope Manager Overview

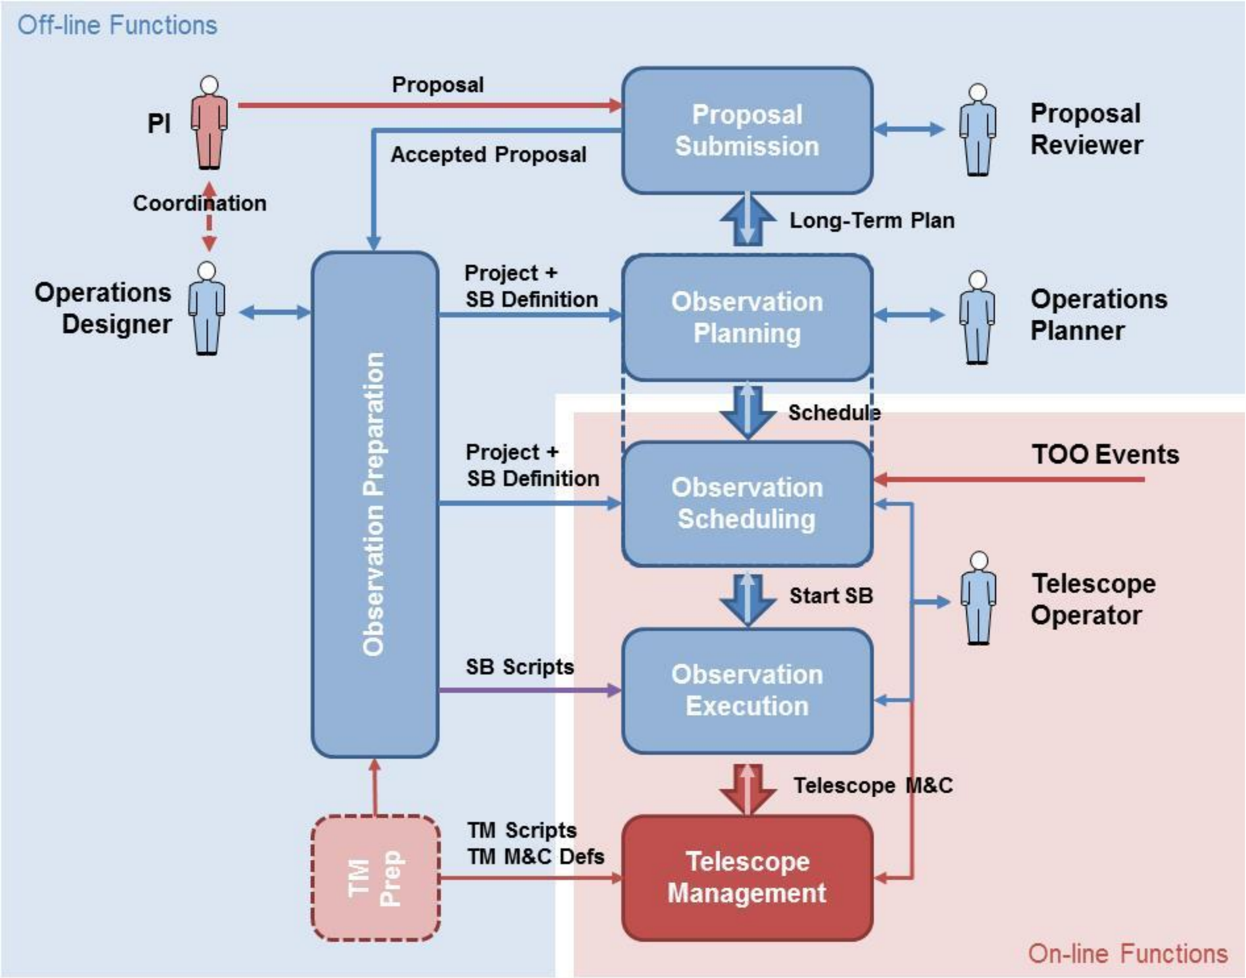

Telescope Management

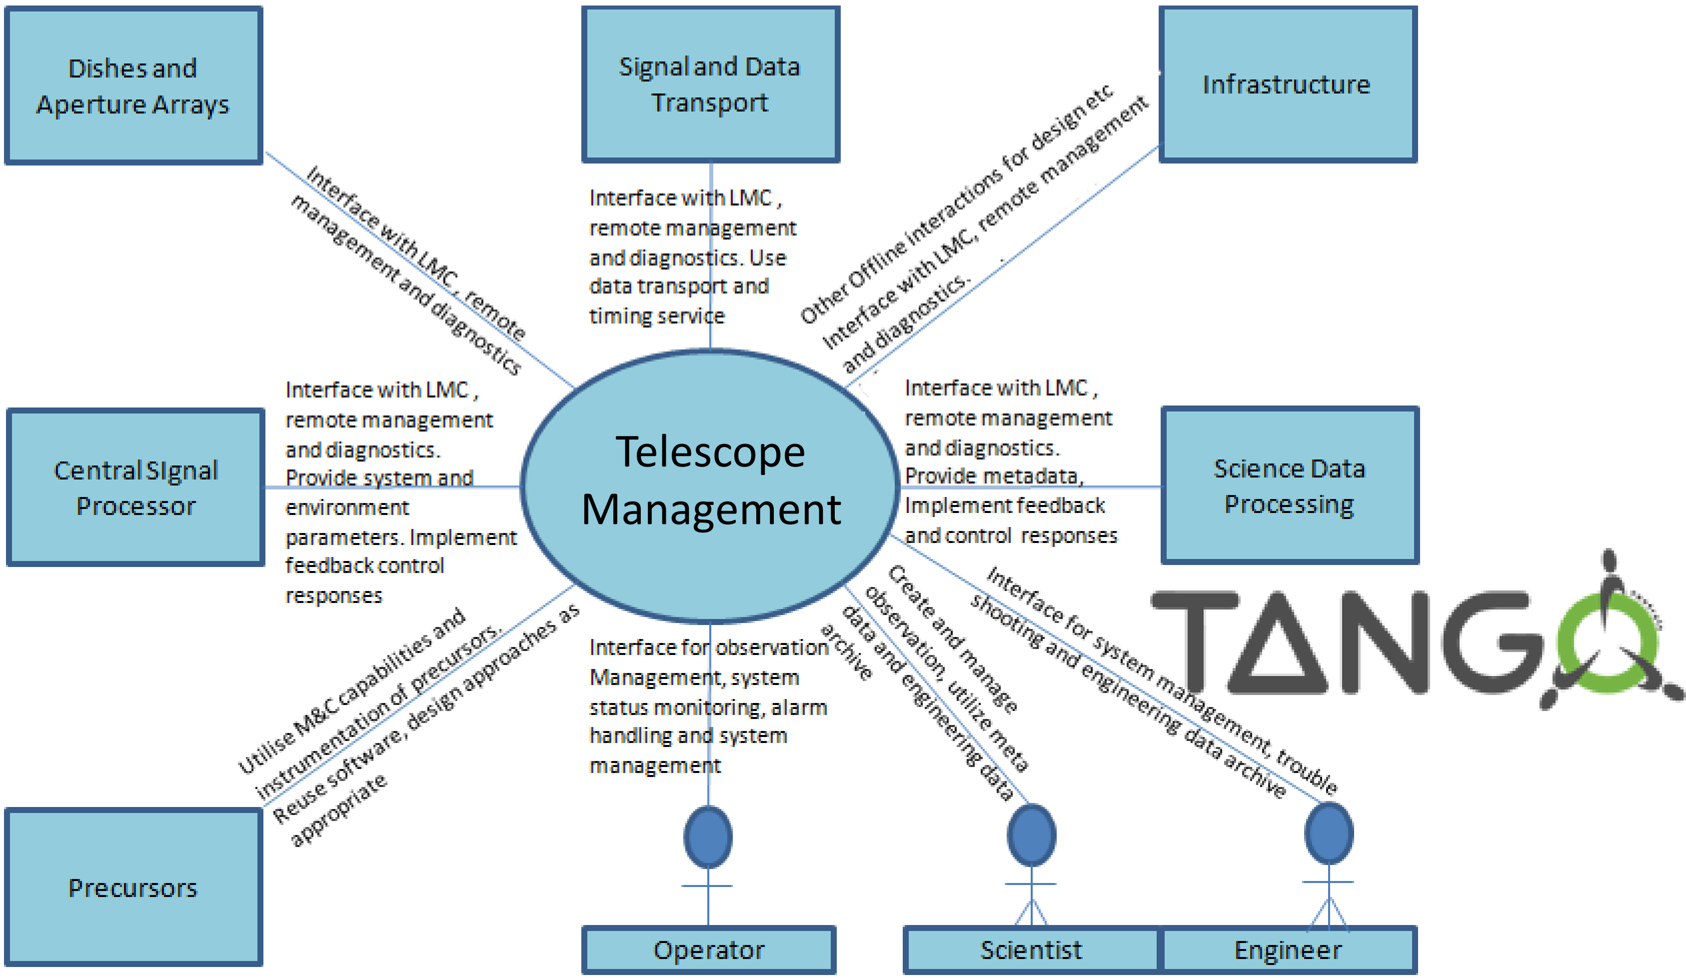

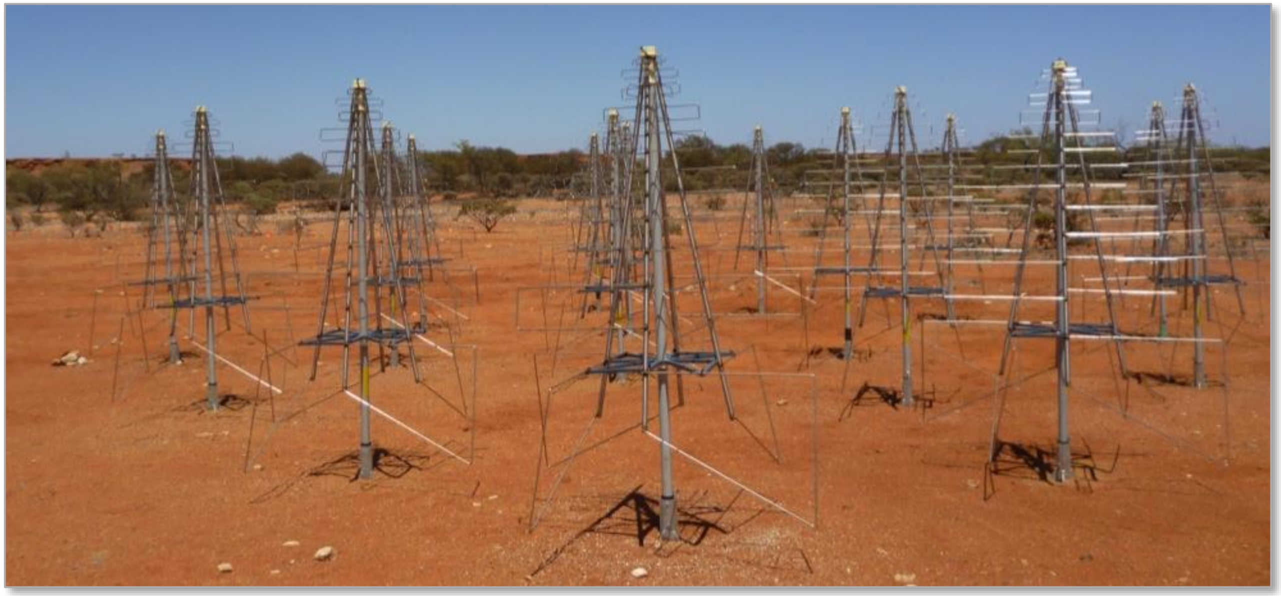

Low Frequency Aperture Array

Technical implementation

Antenna:	Log Periodic
No. of ant.:	131,072 (2^{17})
Ant. Spacing, min:	1.5m (av. ~1.9m)
Station size:	256 antennas ~40m dia.
No. of stations:	512

Signal transport:	Analogue fibre
Processing:	Digital
Sample res.:	8-bit
Grouping:	16 antenna per Tile
Data routing:	**Switched** network

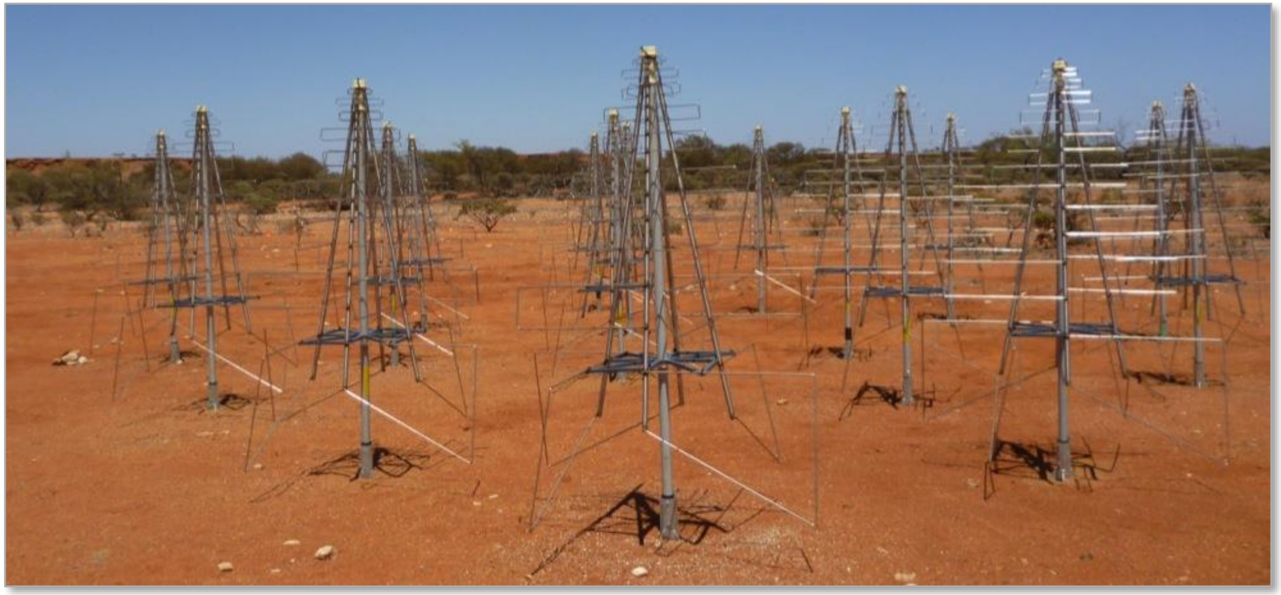

A Station…

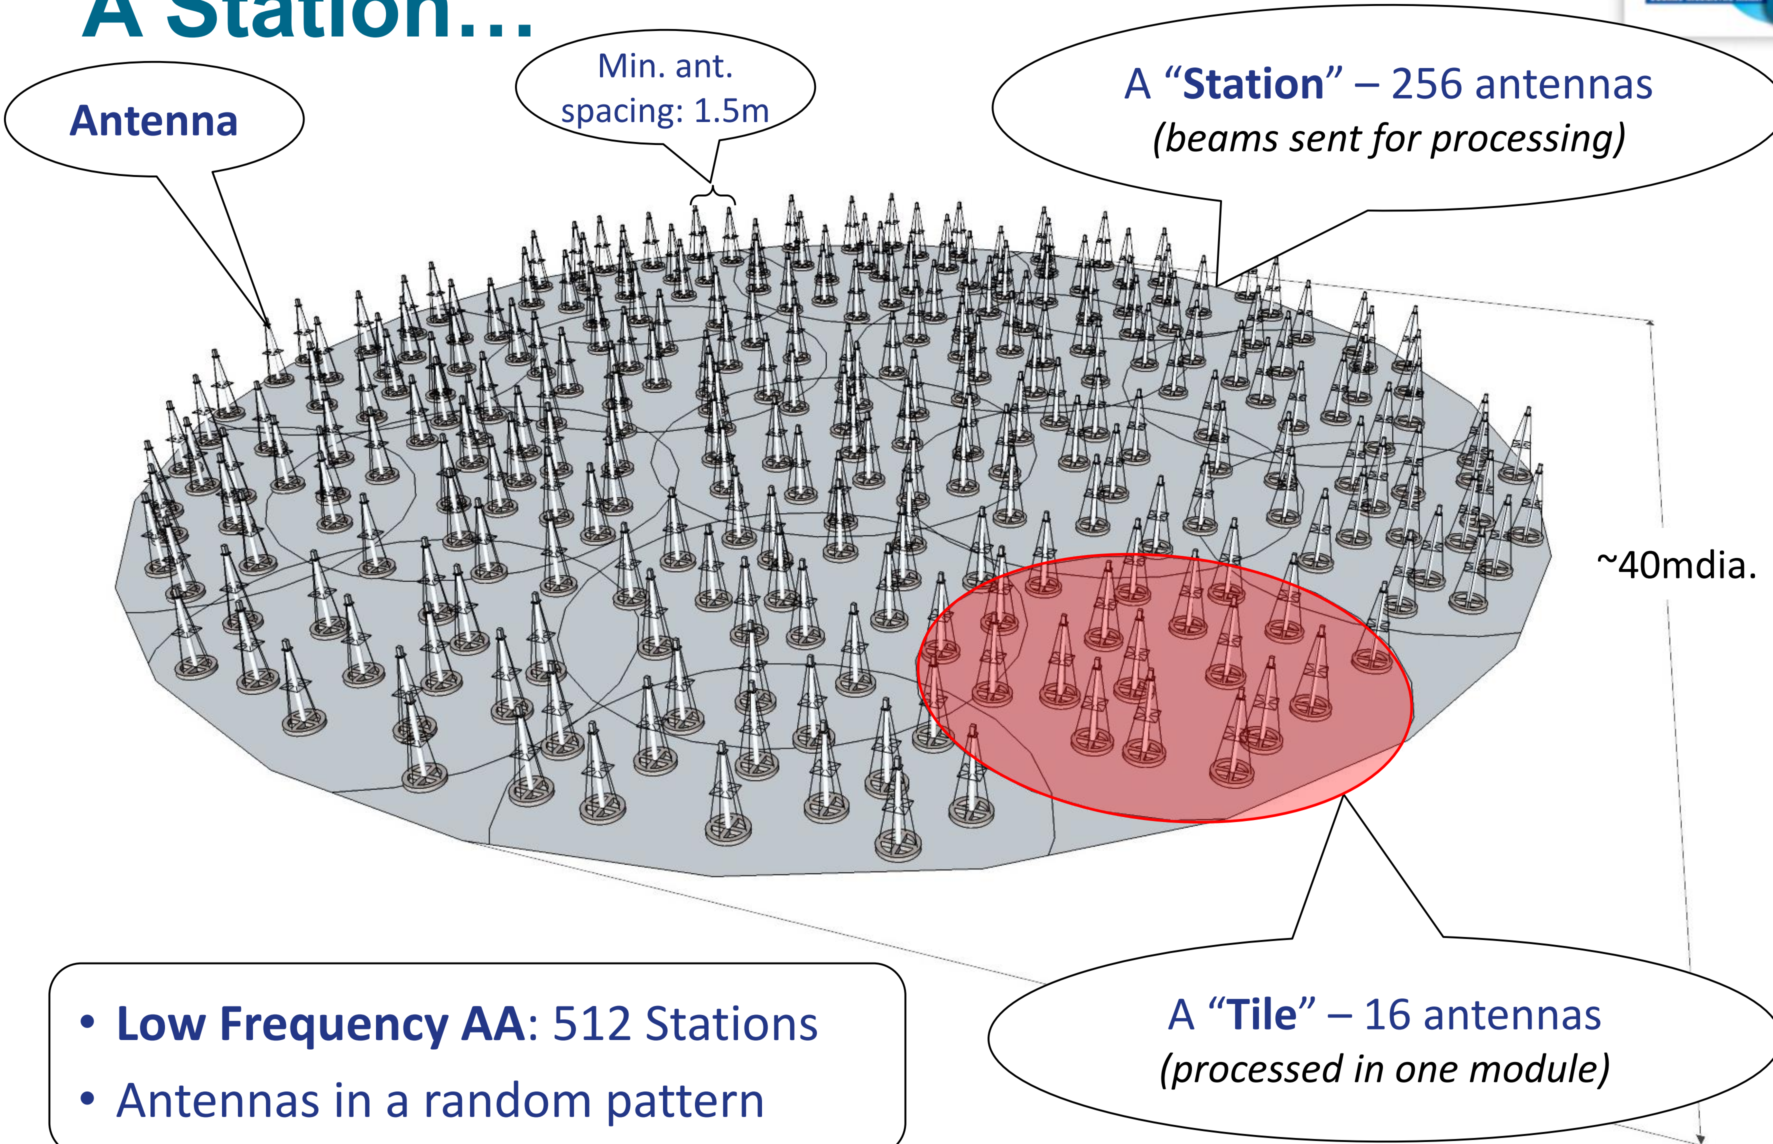

- **Low Frequency AA**: 512 Stations
- Antennas in a random pattern

Low Frequency Aperture Array Overview

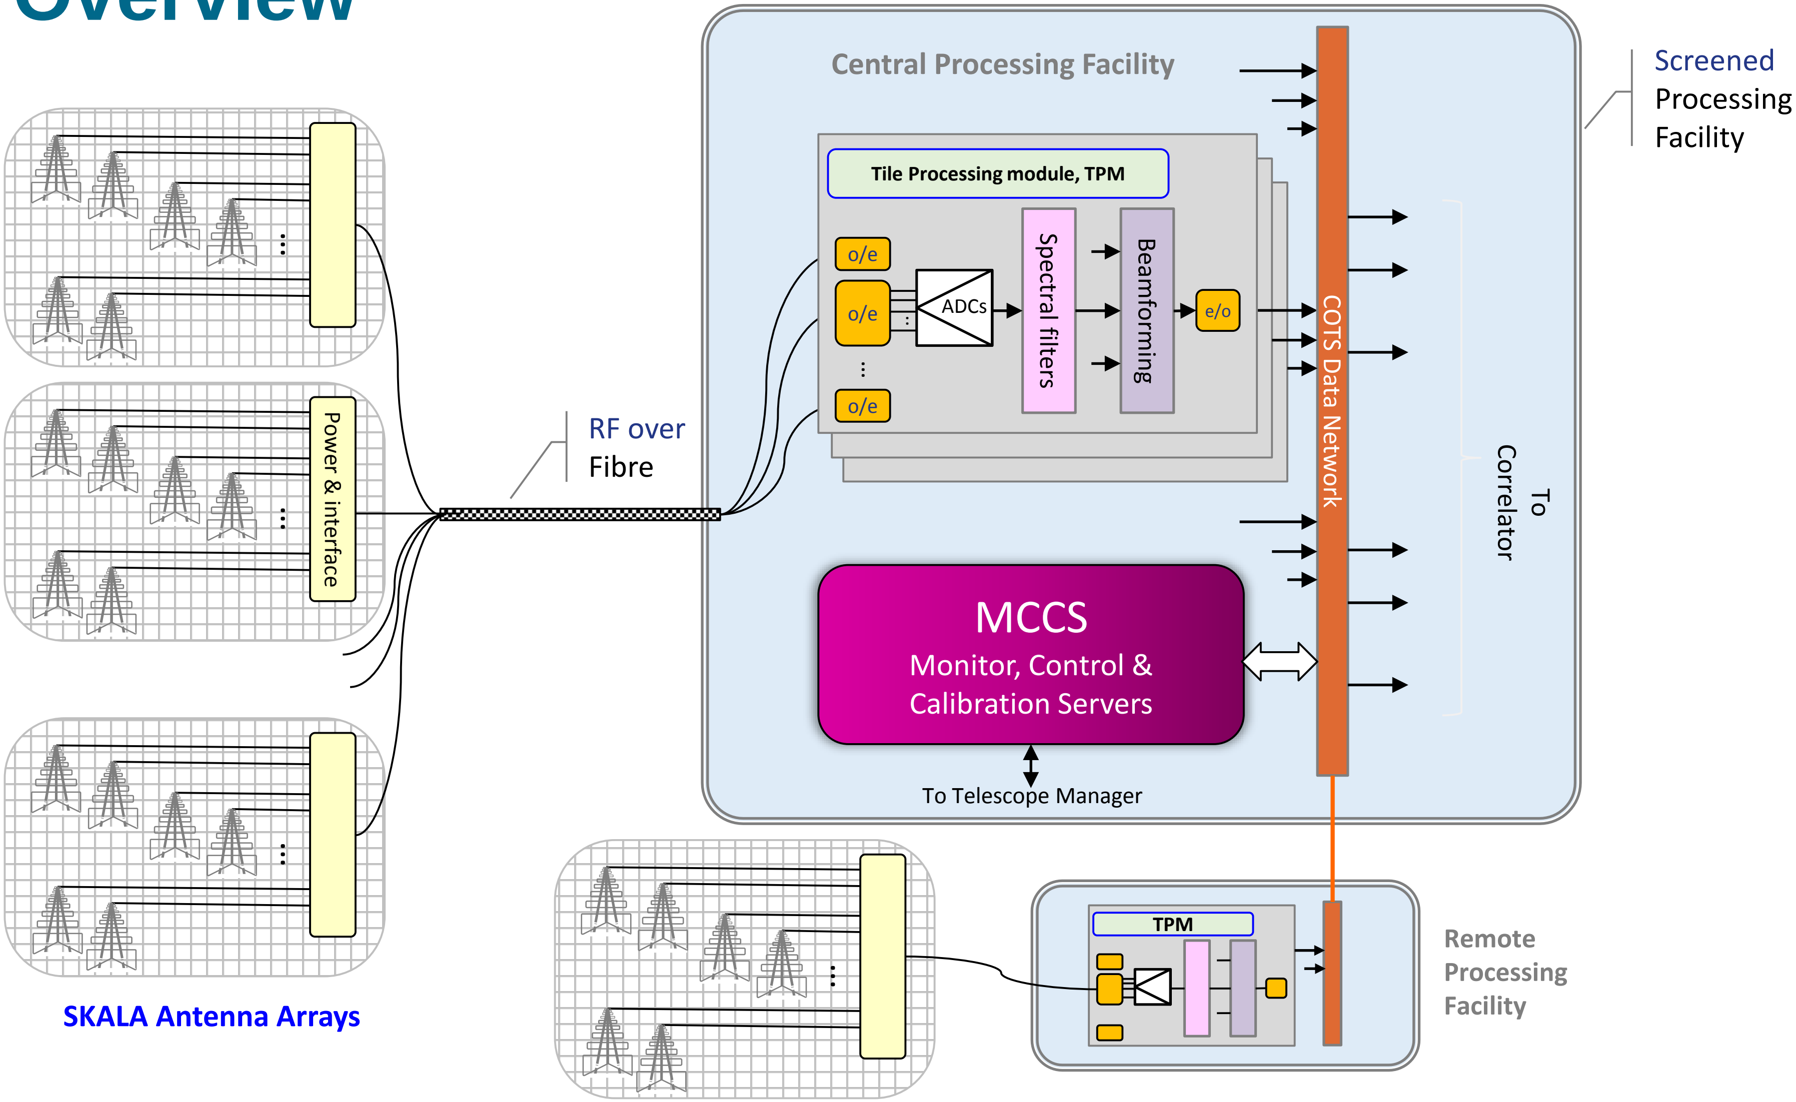

Central Signal Processor (CSP)

Central Signal Processor Overview

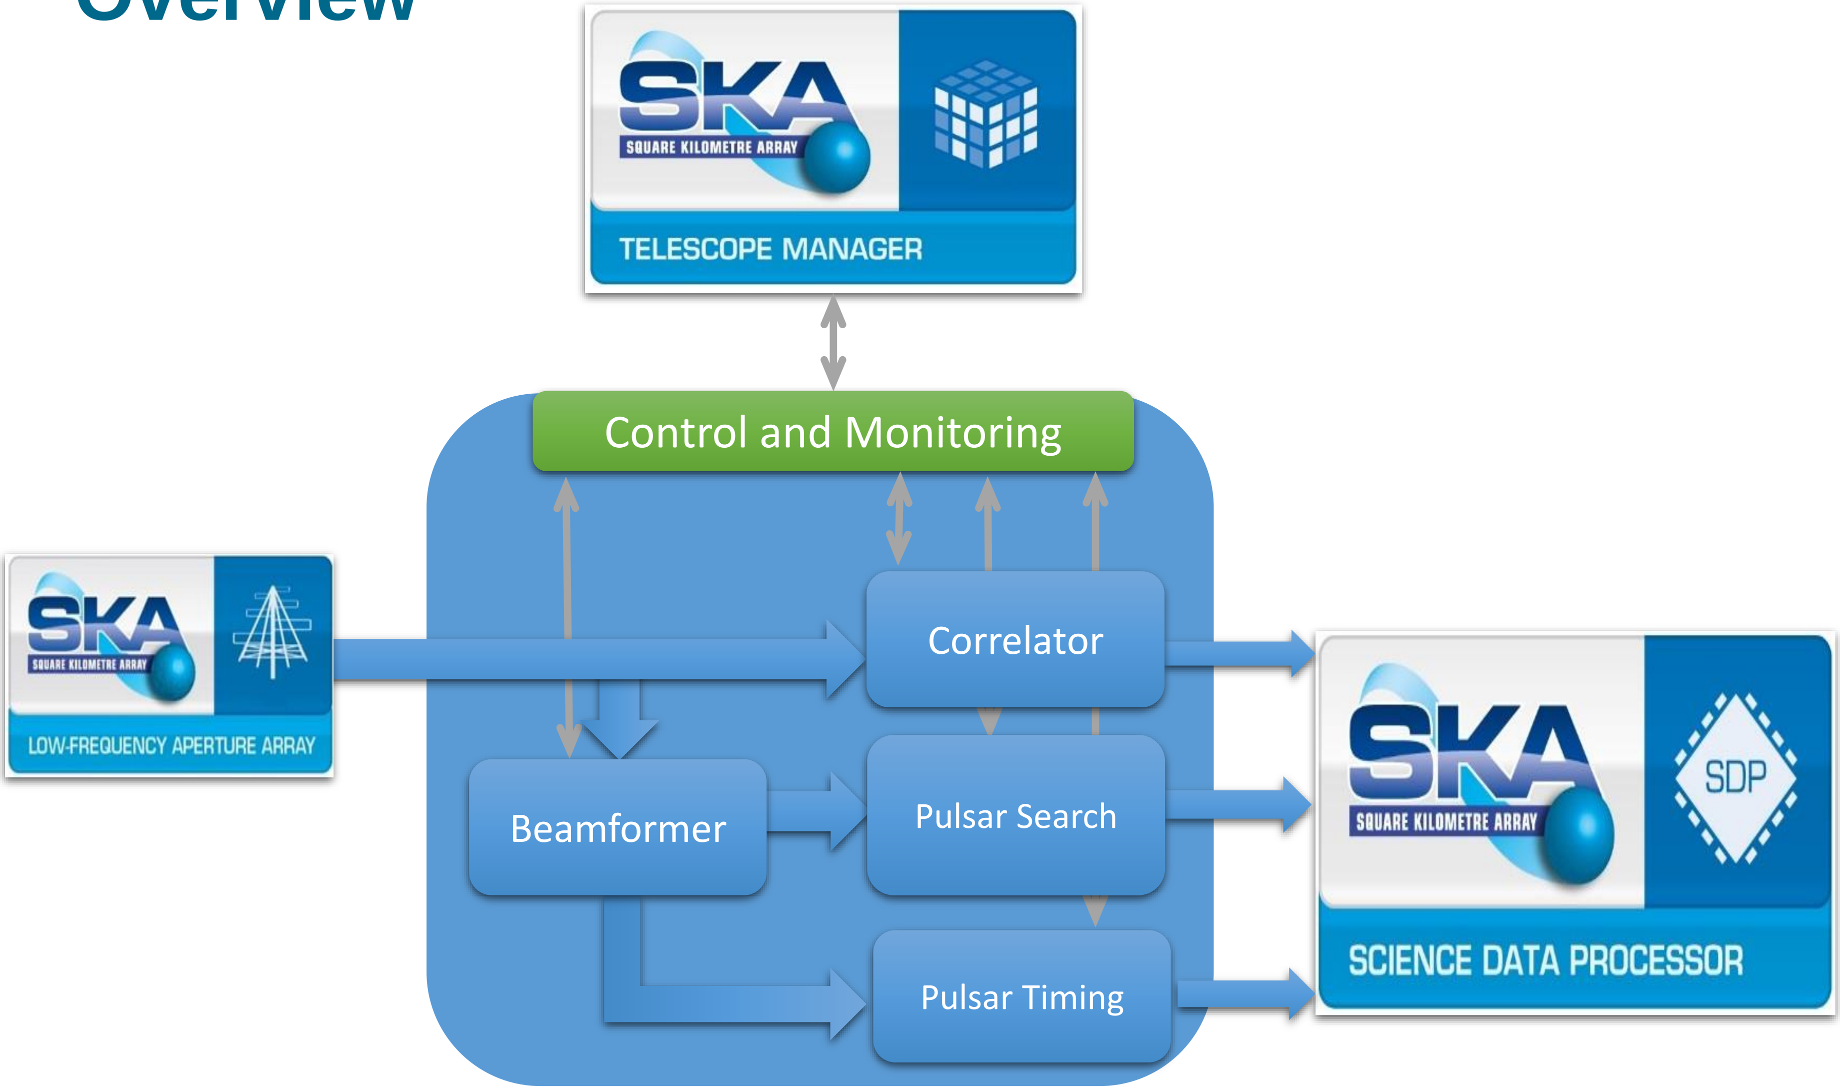

Correlator Beam Former (CBF)

- Correlator:
  - Channelise signal from every dish/aperture array station in to fine frequency channels (65k)
  - Cross-correlate all channels for every pair of dishes/stations
  - Cross-correlations ('visibilities') passed to SDP for imaging
- Beamformer:
  - Forms multiple beams within the dish/station beam
    - 1500 beams for Mid and 750 for Low
  - Passes data to Pulsar Search/Timing engines/VLBI interface
- Very large amounts of real-time processing:
  - N_{corr} ~ $B(N_{dish}.\log_2(N_{ch}) + N_{dish}^2)$ ~ PetaMAC/s
  - N_{BF} ~ $B.N_{dish}.N_{beam}$ ~ few PetaMAC/s
- Based on custom FPGA processing platforms

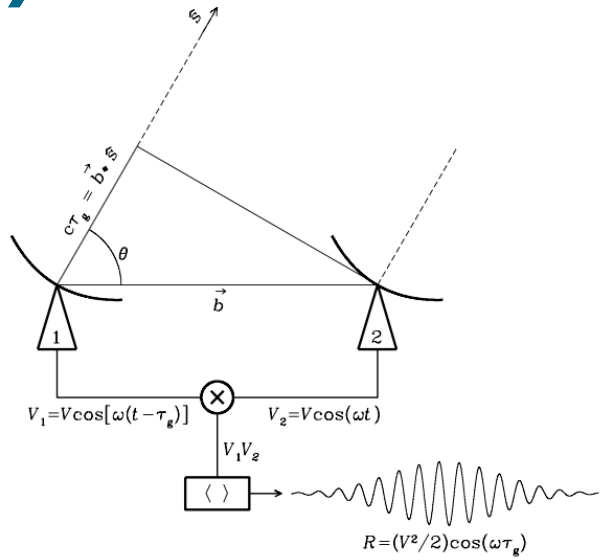

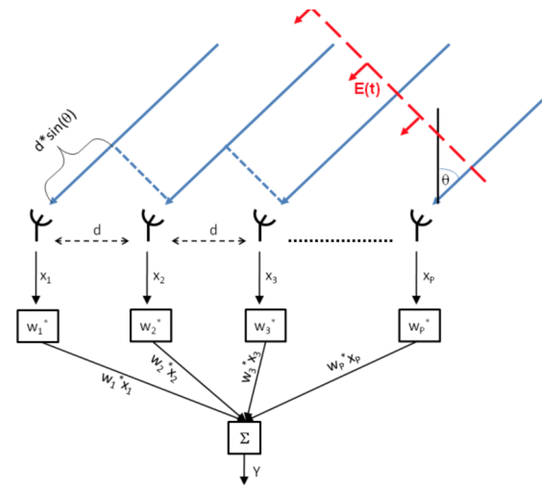

Pulsar Search

- General processing pipeline requires ~50 PFLOPS/sec.
- Baselined heterogeneous design to achieve best combination of hardware & software firmware.
- Two beams per compute node in current design.
- 250 server nodes in Australia and 750 in South Africa
- Dual redundant 10 & 1 gig networks.
- Each Node (1000 in total):
  - Low Power CPUs
  - GPUs
  - FPGA boards
  - 10 Gig inputs
  - > 1 Tbyte RAM &/or SSDs

2+ Gbps in / beam
Real-time dedispersion, 6000 DM trials
2+ Gbps * 6000 out

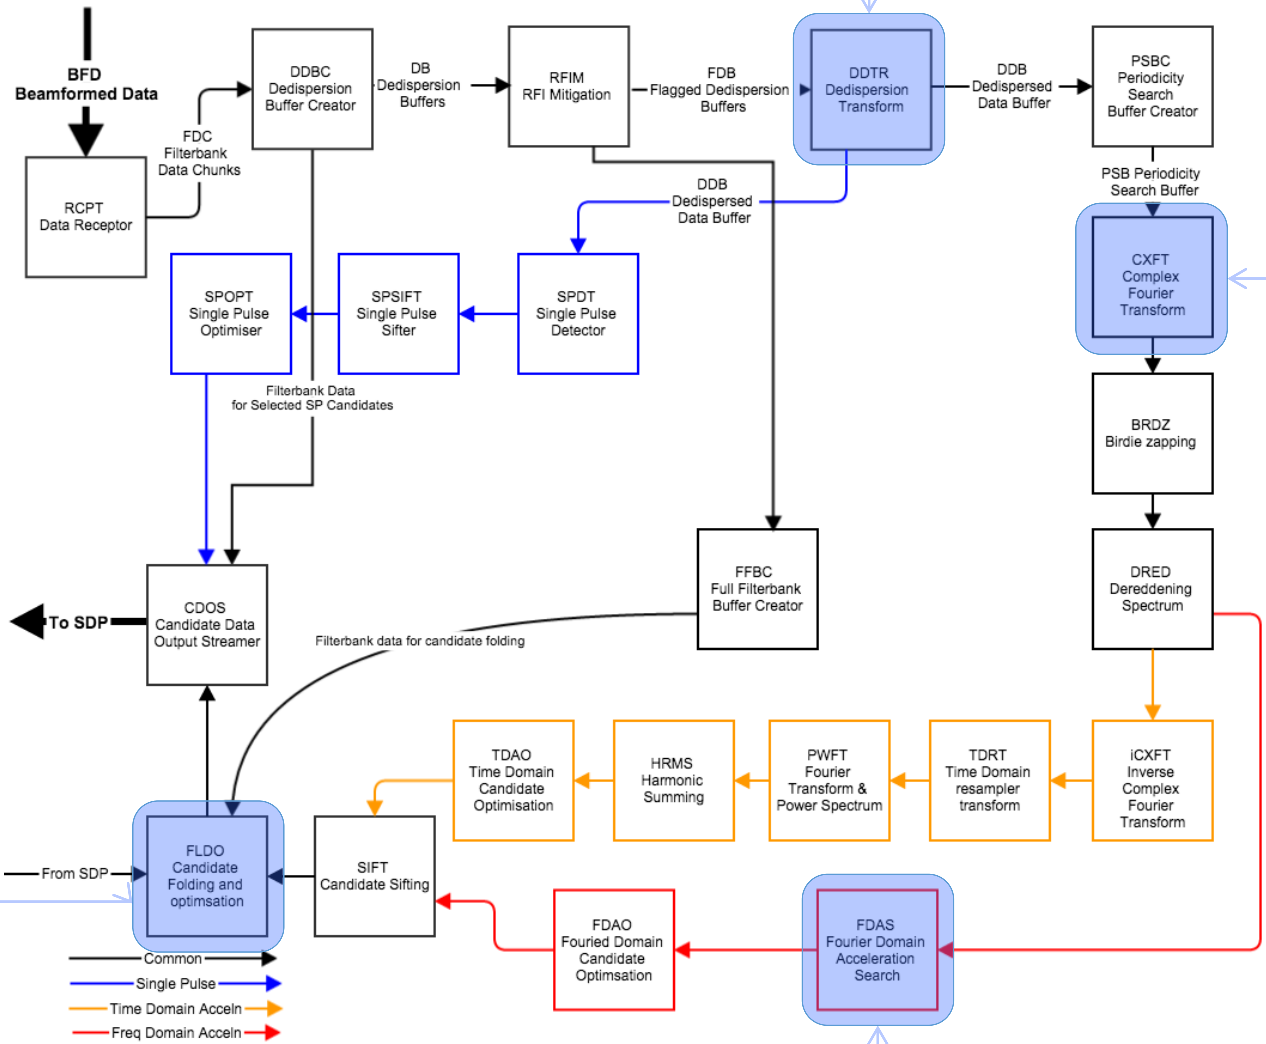

1.5 Gbps in
8Mpt point FFT
10.7 ns per sample
1.5 Gbps out

1000 candidates in
256 kB/folded candidate
380 GB/obs out

1.5 Gbps in
85 * 4 Mpt convolutions
1.5 Gbps * 5 out

Science Data Processor

Science Data Processor Overview

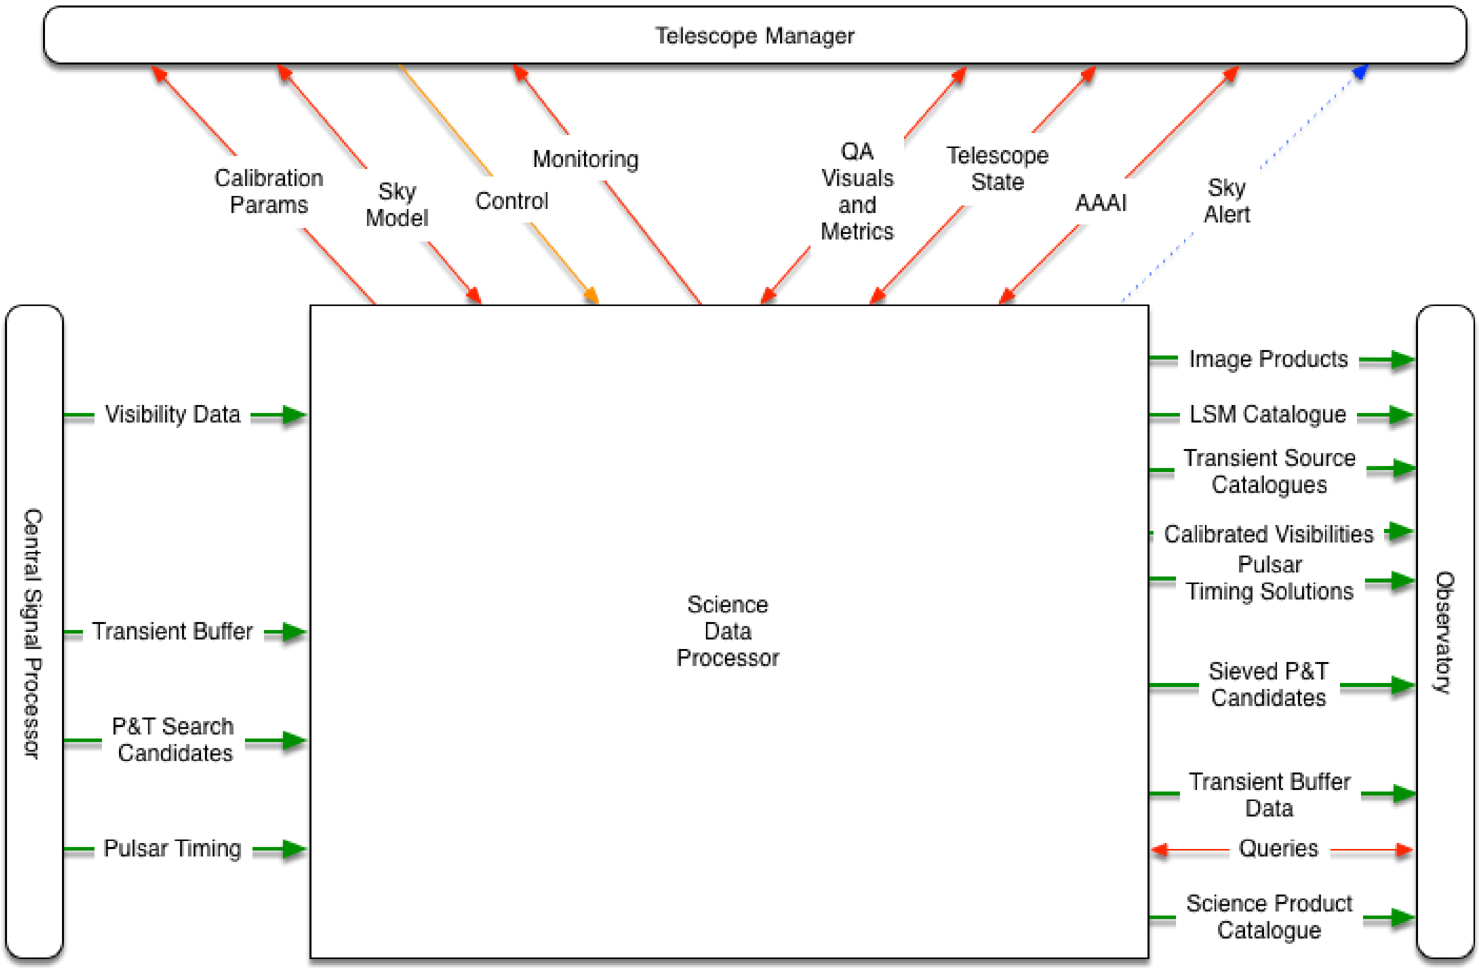

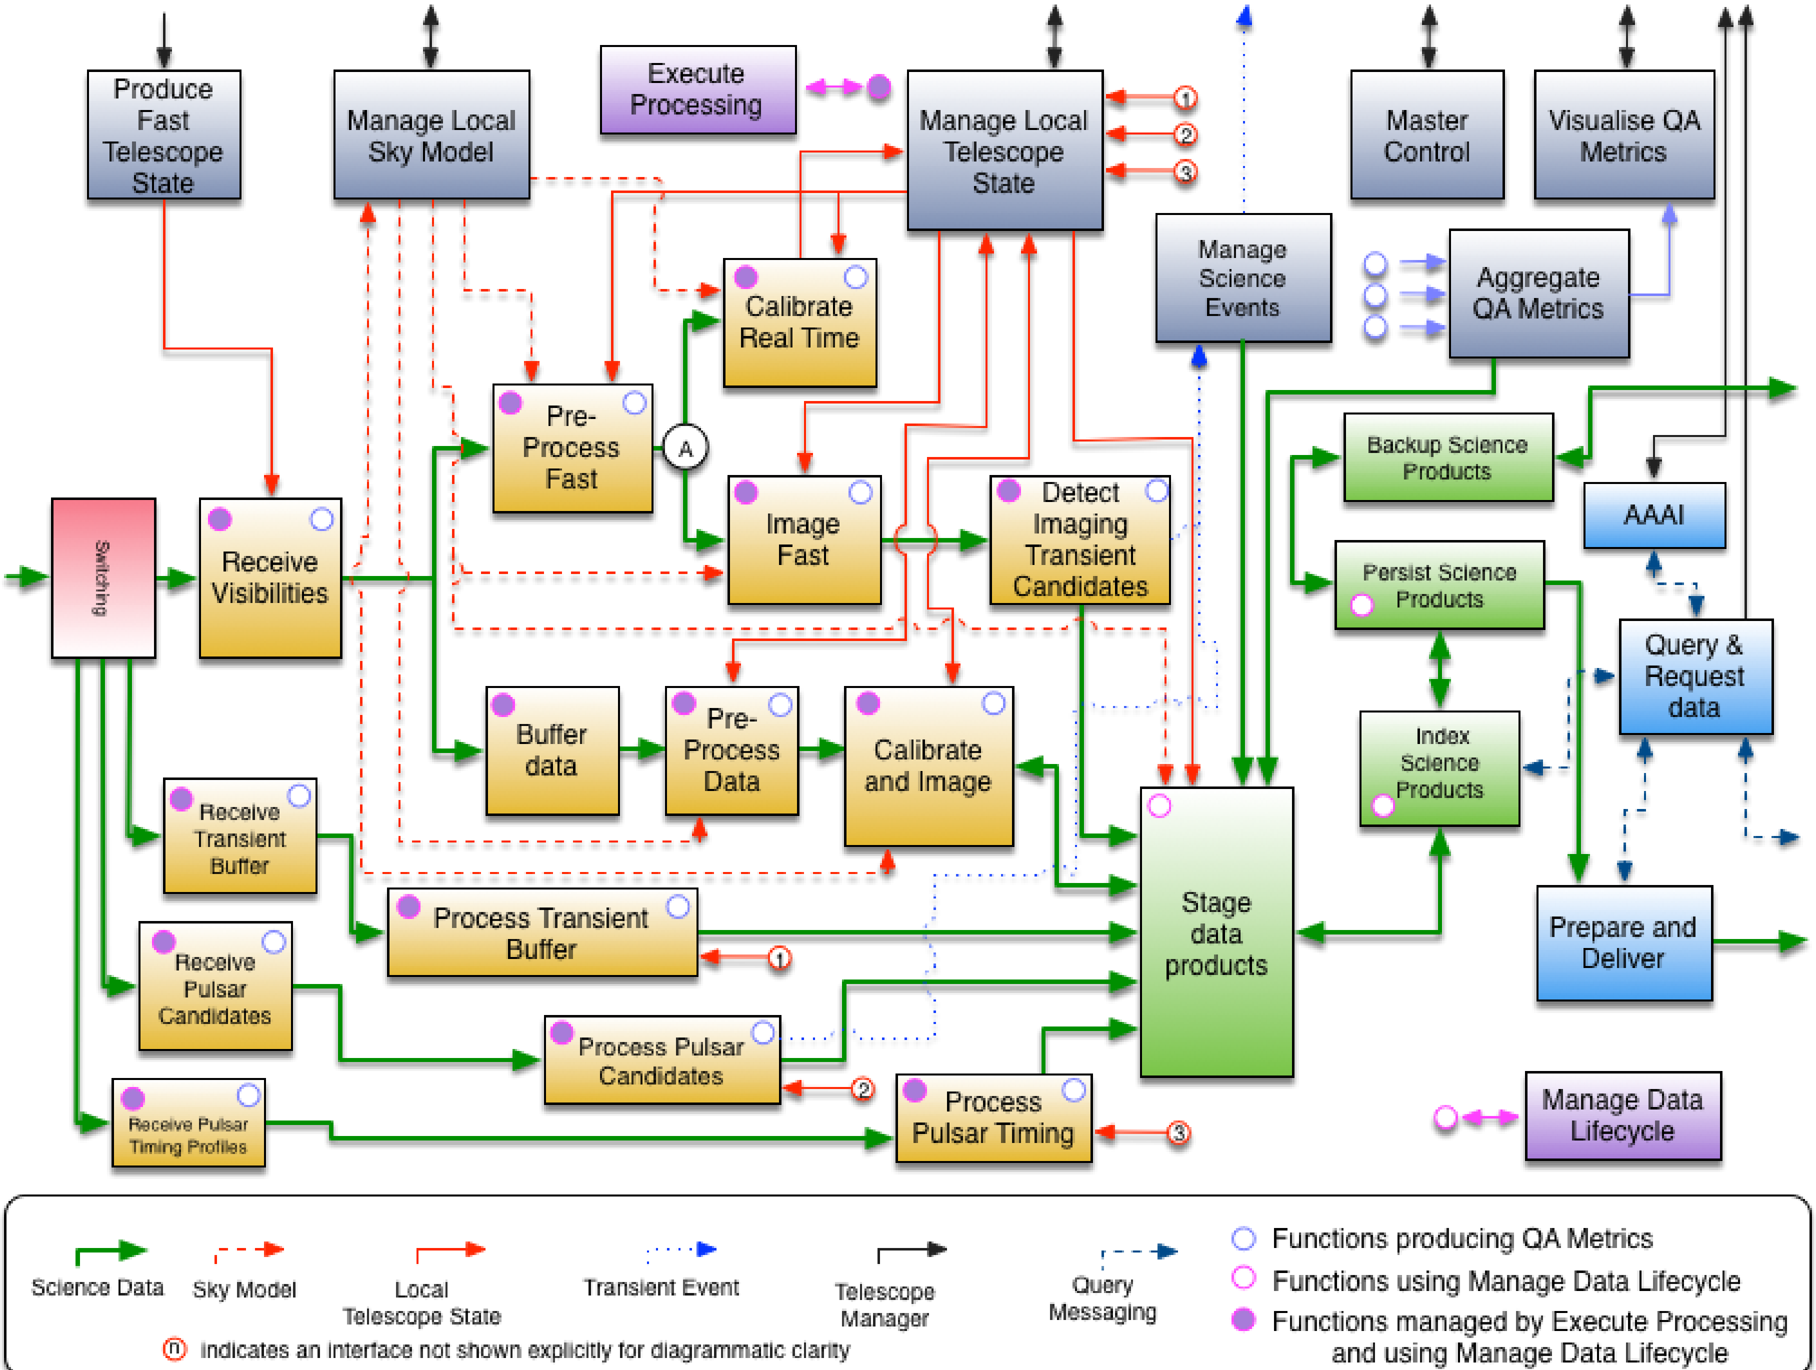
Produce Fast Telescope State
Manage Local Sky Model
Execute Processing
Manage Local Telescope State
Master Control
Visualise QA Metrics
Manage Science Events
Aggregate QA Metrics
Calibrate Real Time
Pre-Process Fast
A
Image Fast
Detect Imaging Transient Candidates
Backup Science Products
AAAI
Persist Science Products
Query & Request data
Switching
Receive Visibilities
Buffer data
Pre-Process Data
Calibrate and Image
Index Science Products
Stage data products
Prepare and Deliver
Receive Transient Buffer
Process Transient Buffer
Receive Pulsar Candidates
Process Pulsar Candidates
Receive Pulsar Timing Profiles
Process Pulsar Timing
Manage Data Lifecycle
1
2
3
Science Data
Sky Model
Local Telescope State
Transient Event
Telescope Manager
Query Messaging
Functions producing QA Metrics
Functions using Manage Data Lifecycle
Functions managed by Execute Processing and using Manage Data Lifecycle
n indicates an interface not shown explicitly for diagrammatic clarity

Graph driven data flow

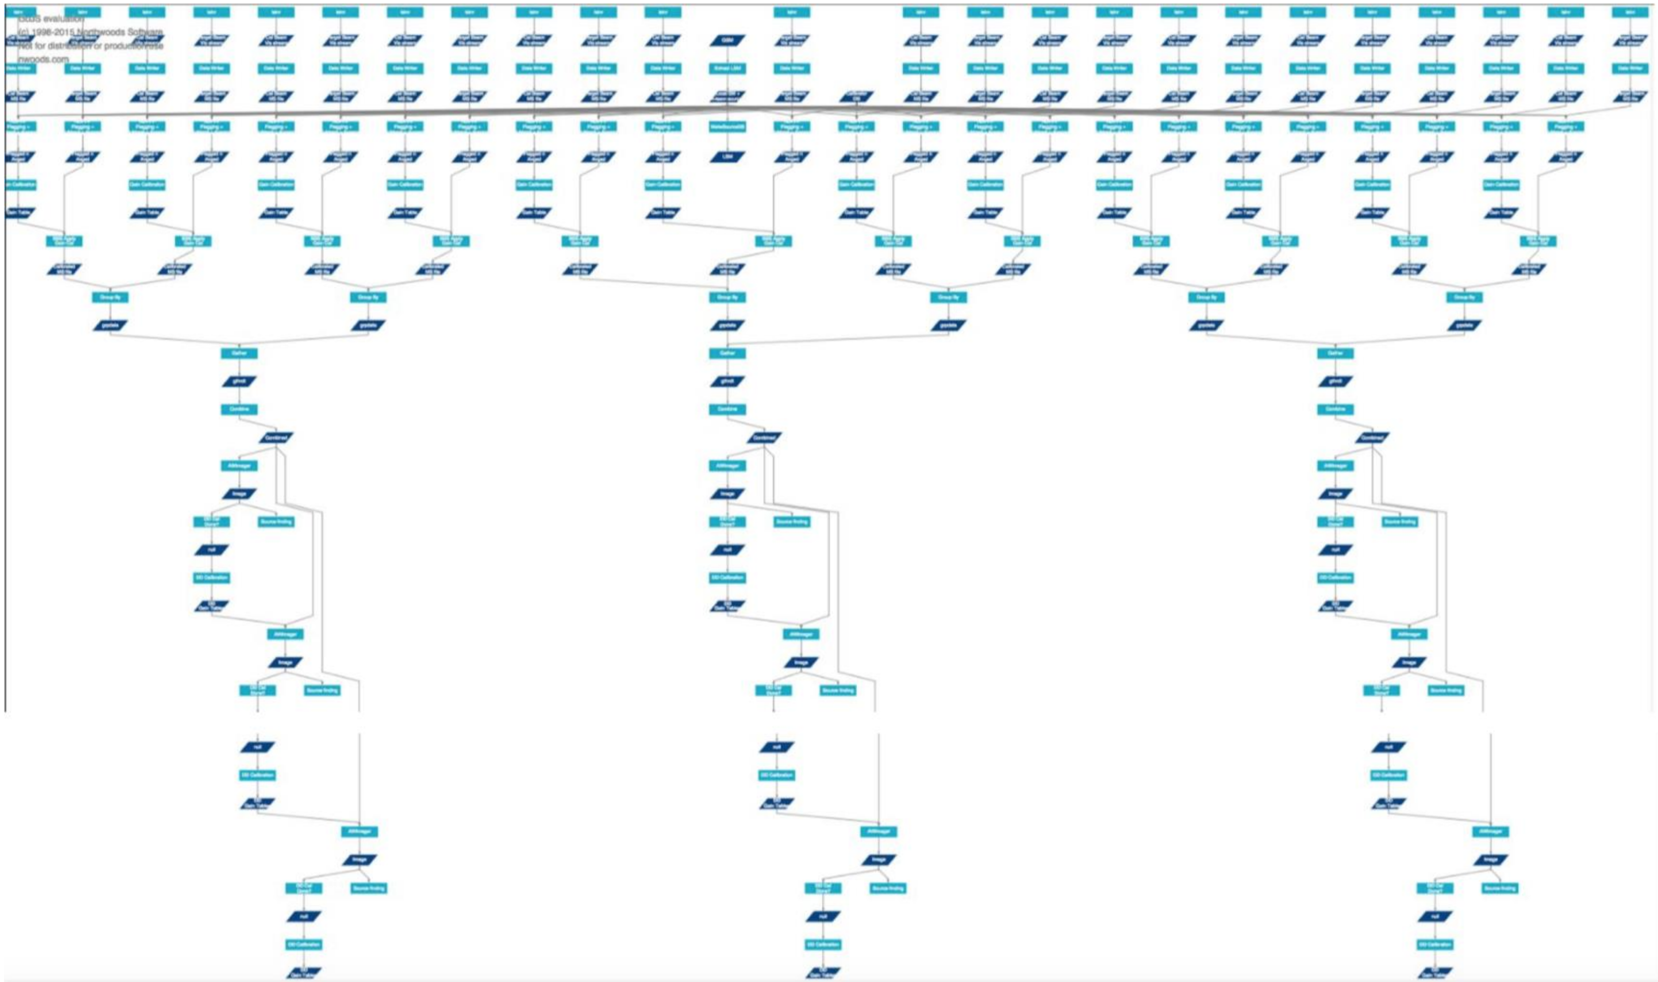

Computing Limitations

- Arithmetic Intensity ρ = Total FLOPS/Total DRAM Bytes
- The principal algorithms required by SDP (gridding and FFT) are typically $\rho \approx 0.5$
- Typical accelerators have an $\rho \approx 5$
  - For example, NVidia Pascal GPU architecture has:
    - Memory bandwidth ≈ 720 GB/sec
    - Floating point bandwidth ≈ 5,000 GFLOPS/sec
- Hence, the computational efficiency ≈ 0.5/5 ≈ 10%
  - So, because of the bandwidth requirements, we have to buy 10 x more computing than a pure HPC system would require.
  - Unless the vendors improve the memory bandwidth…

Computing Requirements

- ~25 PetaFLOPS/sec total sustained
- ~200 PetaByte/s aggregate BW to fast working memory
- ~50 PetaByte fast working storage
- ~1 TeraByte/s sustained write to storage
- ~10 TeraByte/s sustained read from storage
  - ~ 10000 FLOPS/byte read from storage
- Current power cap proposed is ~5MW per site.

Data Management Challenges

- All Top500 HPC systems have been designed for High Performance Computing (by definition).
- There is a new term – High Performance Data Analytics (HPDA) to reflect systems like SKA
- Must ensure the data is available when and where it is needed.
- CPU's must not be idling waiting for data to arrive
  - Data must be in fast cache when it is needed.
- Need a framework that supports this.
  - Looking at a variety of prototypes

Addressing Power

- Need to achieve a FLOPS/Watt 5-10 times better than current greenest computer.
- Need a three pronged approach:

Algorithms	**Hardware**	**Testing**
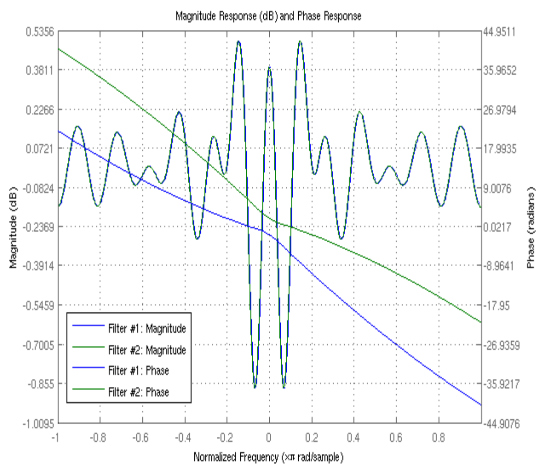	 	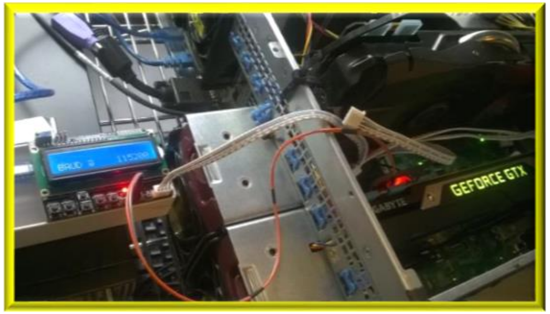 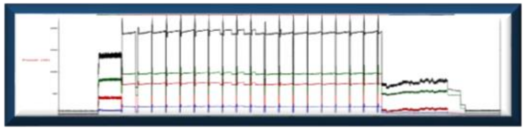
Pursue innovative approaches to cut processing times	Look at accelerators, hosts, networks and storage.	Using real algorithms and fully instrumented systems

Software

- Budget of ~€90M on manpower for software development across the whole telescope.
  - Need professional practices for development, testing, integration and deployment.
  - Need to unify the processes across the world-wide team of developers.
  - Need world-leading expertise in a number of areas.
- Delivered system will not be static
  - SDP hardware and software will be updated periodically.
  - Key input for development will be the scientific and software community through the regional centres.

SKA Regional Centres

Regional Centre Overview

- In April 2016 the SKA Board agreed the principal of regional centres
- Modelled on the LHC Tiering system
- Not part of SKA, or funded by SKA, but:
  - Essential to generate science
  - Coordinated with assistance from SKAO and accredited with SKAO
- Principle functions
  - Take data products generated by SDP and turn them into science
  - Support regional astronomers with their data processing.
  - Act as a centre for domain expertise.

Regional Centre Concept:

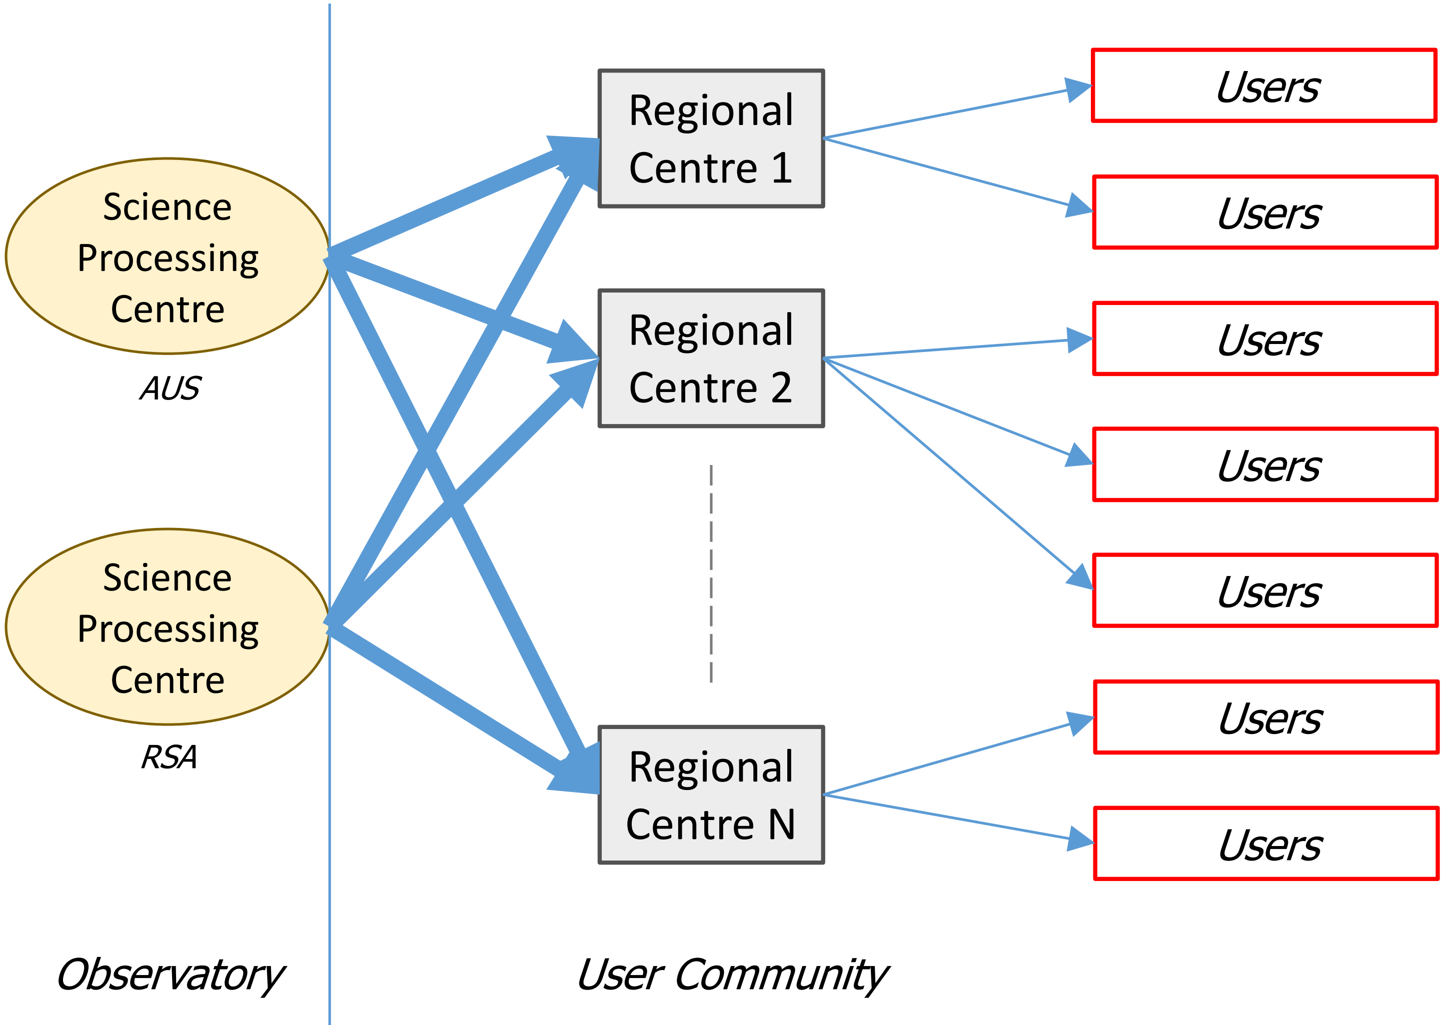

Regional Centre Network

- 10 year IRU per 100Gbit circuit 2024-2033 (2015 est.)
- Guesstimate of Regional Centre locations

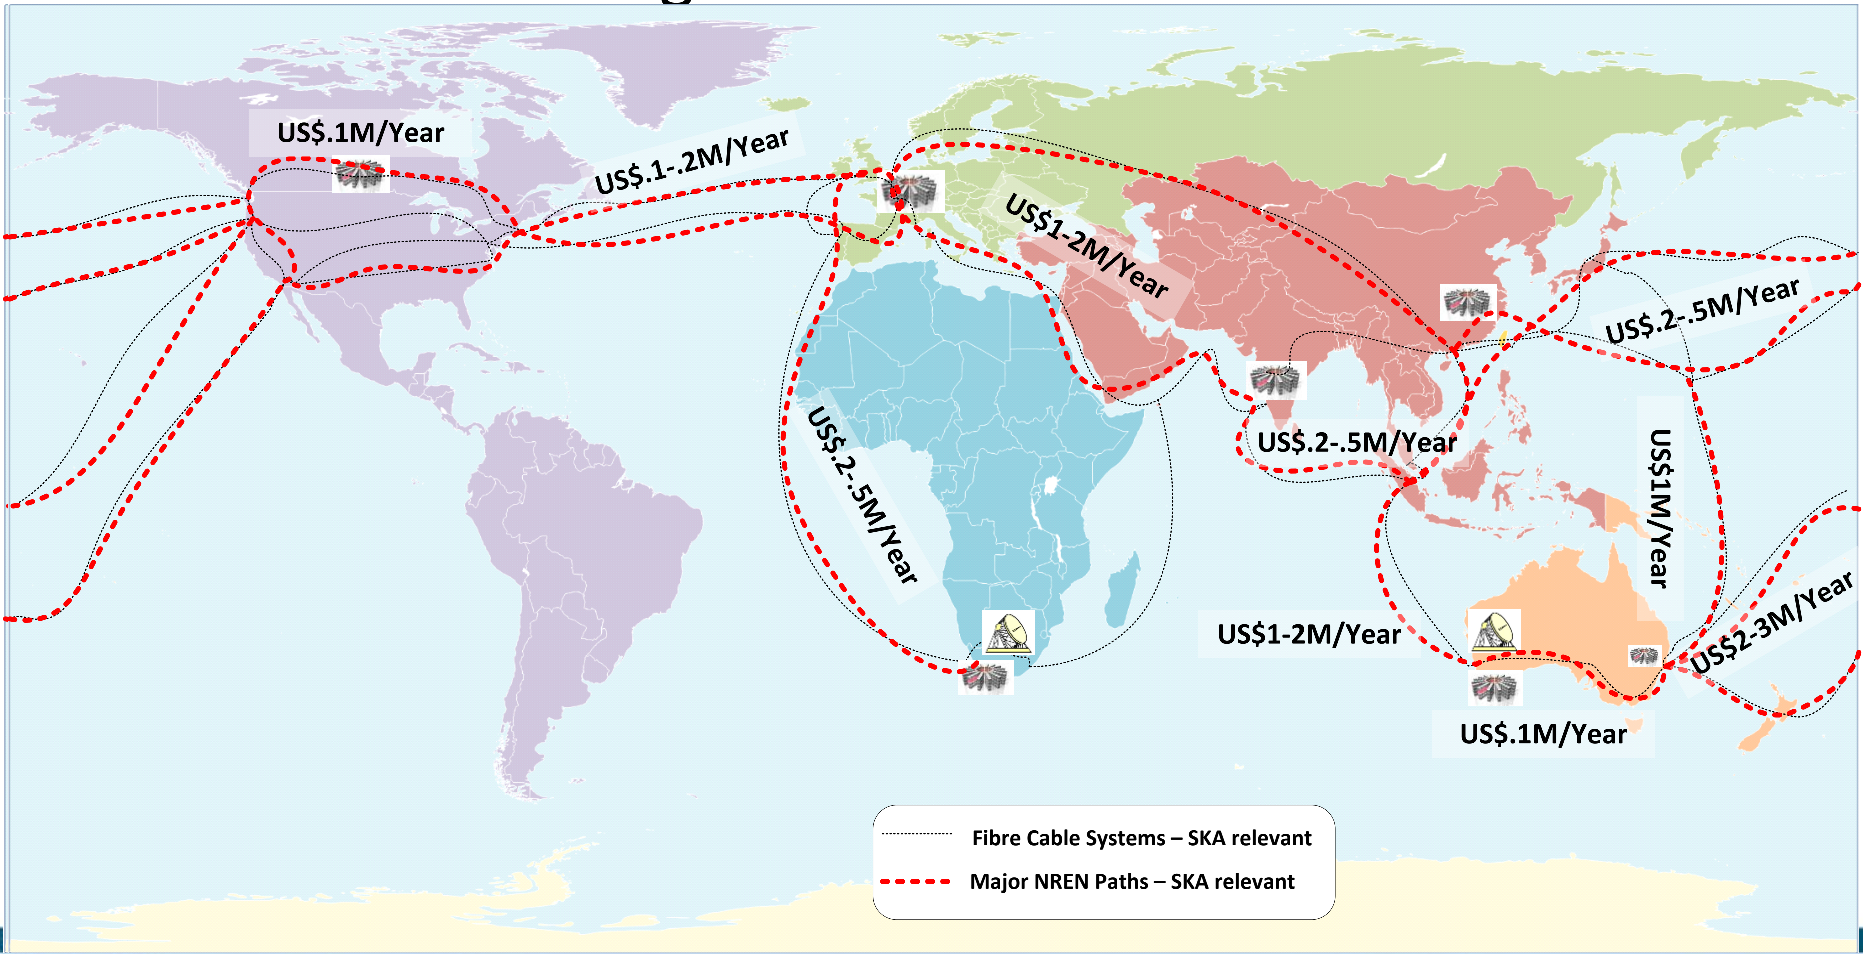

What will Regional Centres do?

- Provide a nexus for resources
  - Scientific expertise
  - Software expertise
  - Access to computing resources - but direct ownership of is not a requirement
- Provide support for scientific development
  - Provide future, subject-specific pipelines
  - Contribute to common efforts in visualization, stacking, co-adding?
- Provide access to data
  - Ensure security and adherence to SKA data policies
  - Play a defined role in hosting and distributing archives
- Provide local (time zone) user support, proposal access, information, training and outreach activities
- Liaison with SKA Observatory and NRENs
  - Need to ensure sufficient and affordable network capacity is procured and provisioned in a timely fashion

Conclusions

- SKA will be the world's primary radio telescope in the metre and centimeter bands.
- It is a huge computational and software challenge
- Traditional HPC is not a good match because the problem is bandwidth dominated.
  - SKA is seen as a key programme in global IT development
  - Showcases a major development area of High Performance Data Analysis (HPDA).
- Software complexity is also beyond what has been achieved in astronomy previously.
  - Quality is paramount and good processes essential

Questions?