Data Processing for the Square Kilometre Array Telescope

Streaming 2015 Workshop – Indianapolis

27th October 2015

Bojan Nikolic

Astrophysics Group, Cavendish Lab

University of Cambridge

Email: b.nikolic@mrao.cam.ac.uk

Project Engineer & Architecture Task Leader

SKA Science Data Processor Consortium

Not covered in this talk

Not covered:

- Science Objectives of the SKA
- Timelines, project programmatic and organisational structure
- Design of the receptors (aperture arrays & dishes) and their associated digital electronics/communications equipment

Also quite a few spare slides on topics covered.

Find me later if you'd like information on these

THE SQUARE KILOMETRE ARRAY TELESCOPE

What is the Square Kilometre Array (SKA) ?

SCIENCE DATA PROCESSOR

Mid frequency dishes and dense aperture array concept

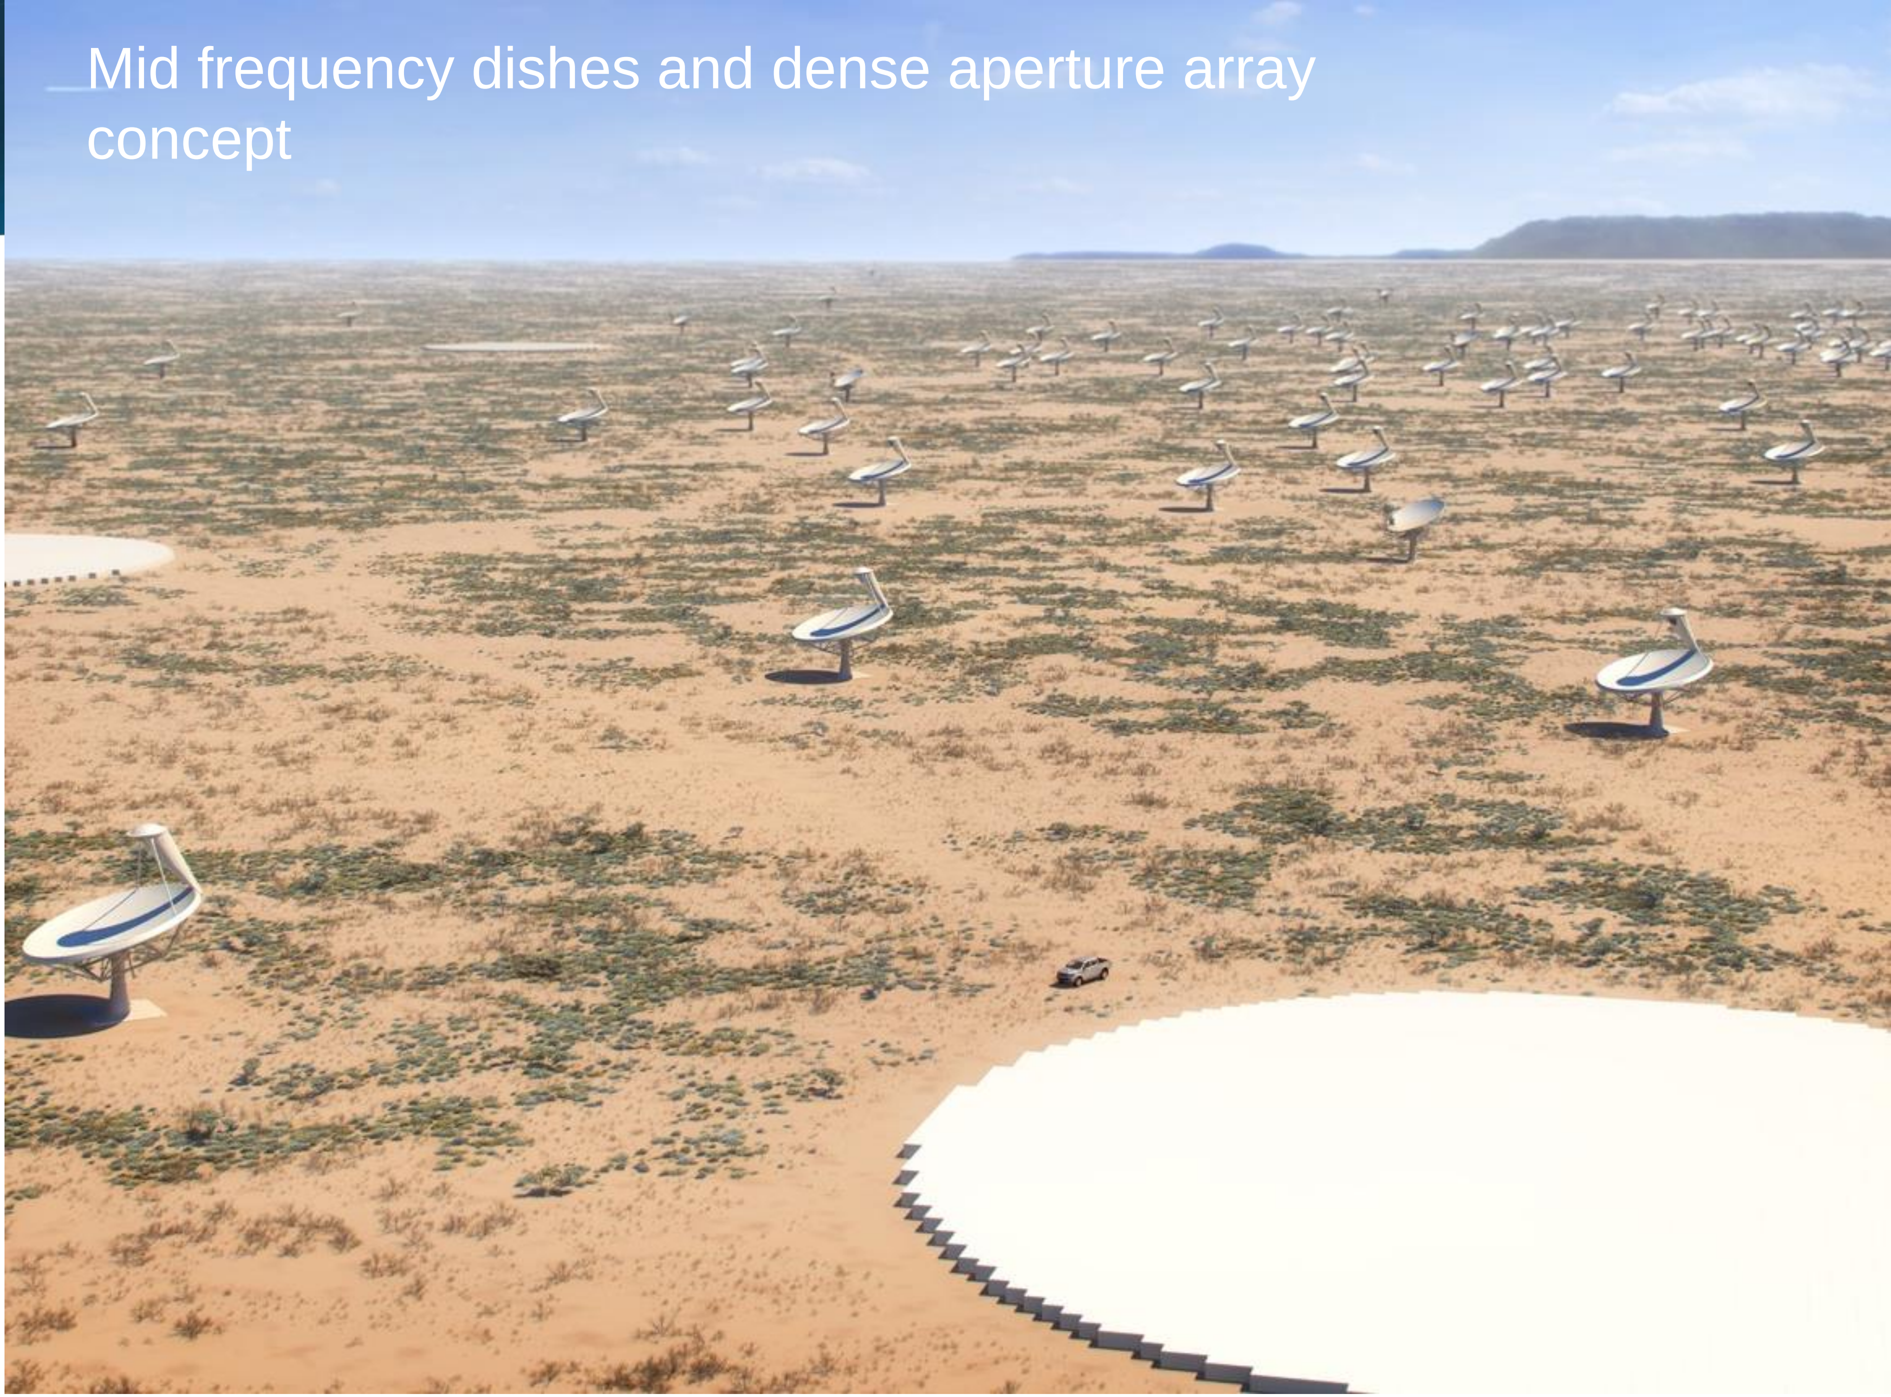

Low frequency array and the survey telescope concept

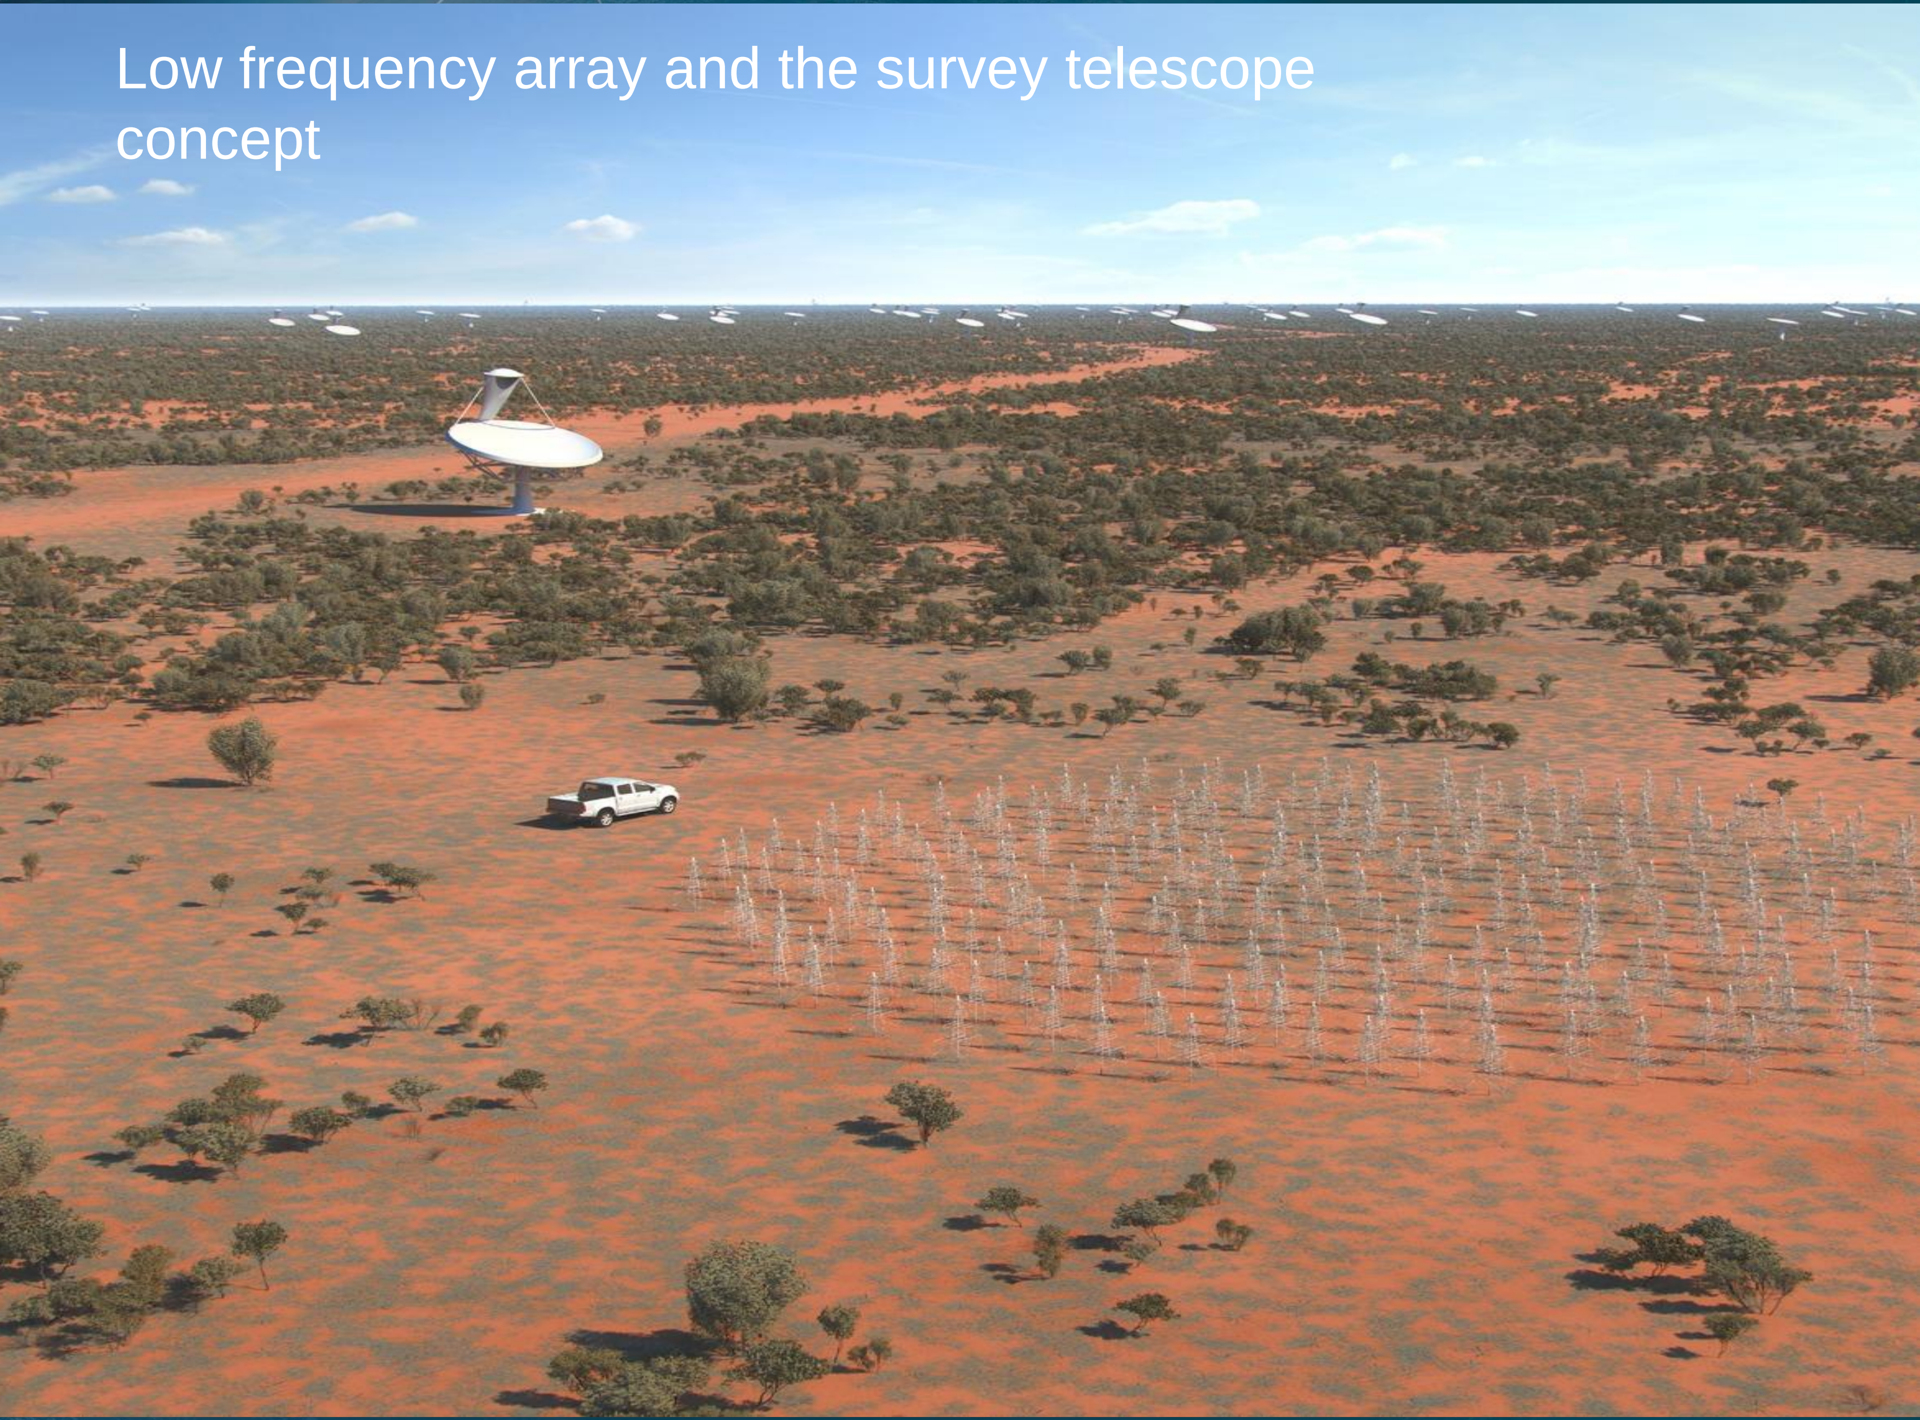

Phased Aperture array for 40 – 650 MHz

Phased Aperture array: 3 million antennas

Scientific Context – a partner to ALMA, EELT, JWST

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

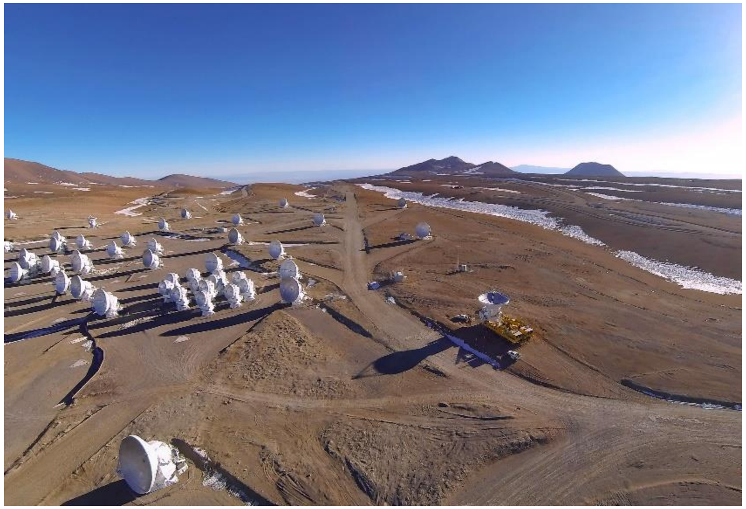

Credit:A. Marinkovic/XCam/ALMA(ESO/NAOJ/NRAO)

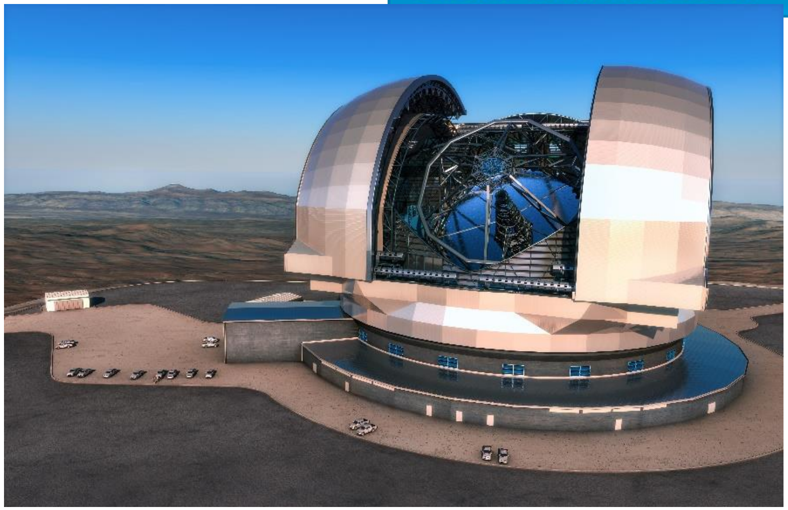

Credit:ESO/L. Calçada (**artists impression**)

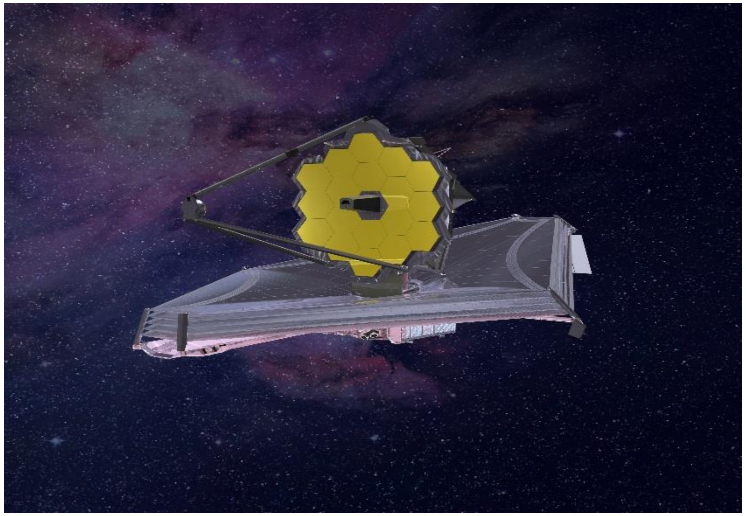

Credit: Northrop Grumman (**artists impression**)

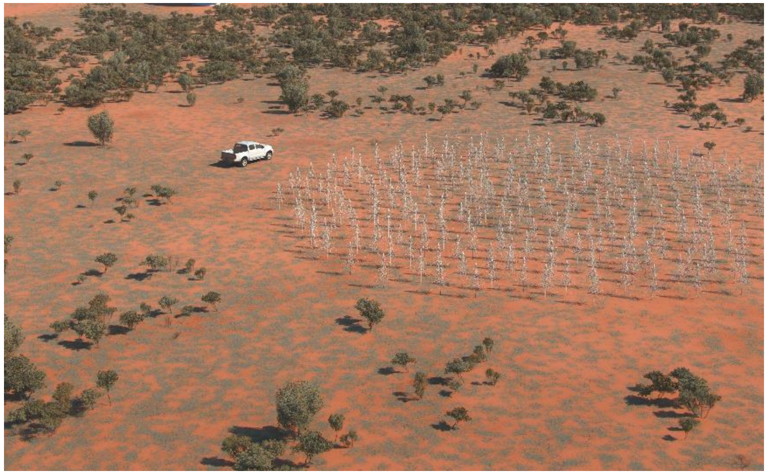

Credit: SKA Organisation (**artists impression**)

Scientific Context – a partner to ALMA, EELT, JWST

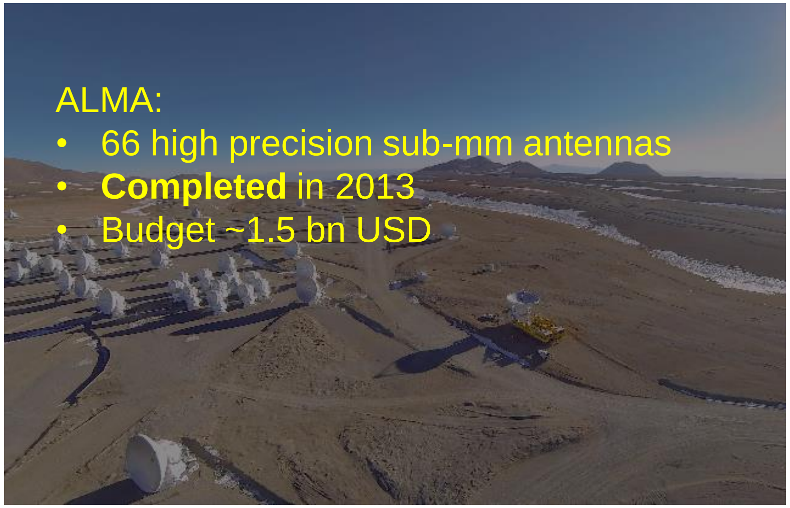

Credit:A. Marinkovic/XCam/ALMA(ESO/NAOJ/NRAO)

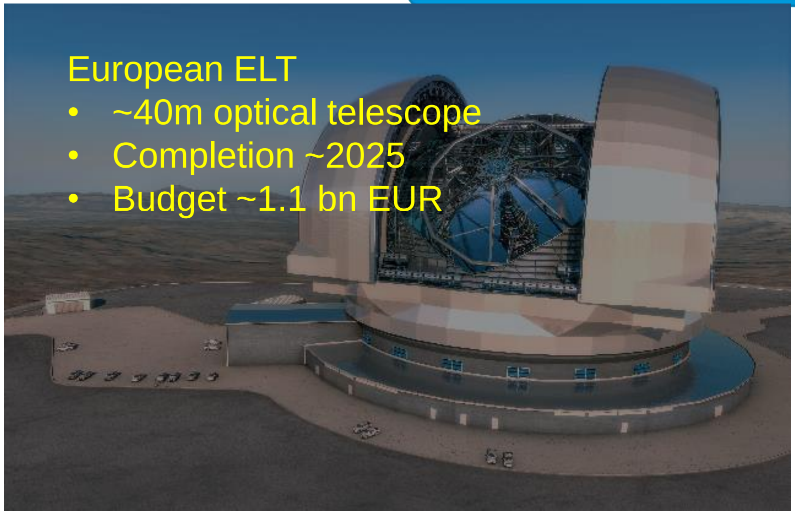

Credit:ESO/L. Calçada (**artists impression**)

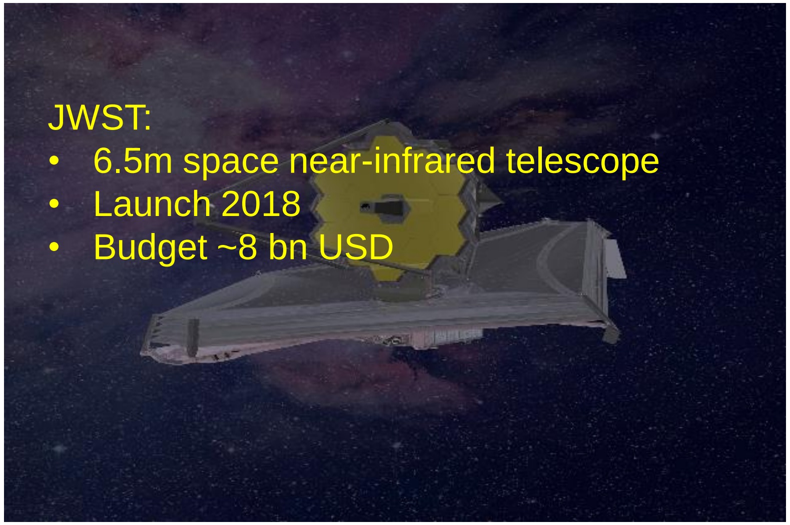

Credit: Northrop Grumman (**artists impression**)

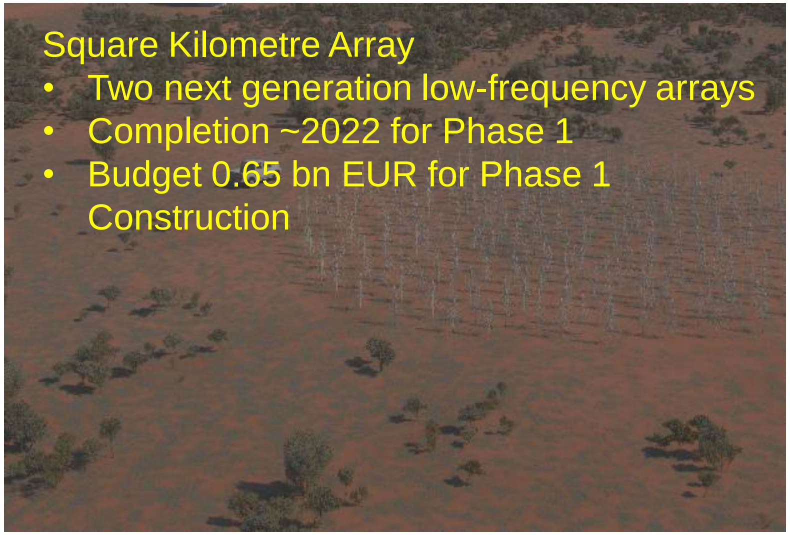

Credit: SKA Organisation (**artists impression**)

What will the Square Kilometre Array (SKA) be?

SCIENCE DATA PROCESSOR

Radio Telescope	Makes Images of the Sky at radio (5m-3cm) wavelengths
	~100 more sensitive than current telescopes
	Complements ALMA, JWST (successor to Hubble), and E-ELT
Currently in Design	Construction begins 2018
	Full operations expected at end 2022
	Significant funds already committed by participating countries
Major Engineering Project	Two remote desert sites
	>100k receiving elements
Major ICT Project	Subject of this talk!

SCIENCE DATA PROCESSOR CONTEXT

The SKA Observatory – Phase 1 : “SKA$_1$”

+ People, buildings, roads, ground works, communications, electrical power, maintenance, transportation, catering at desert sites….

Data Processing: general computing

Digital Signal Processing: FPGA + (maybe) ASICS & GPUs

Receptors: Aperture Arrays and Dishes

SKA Context Diagram

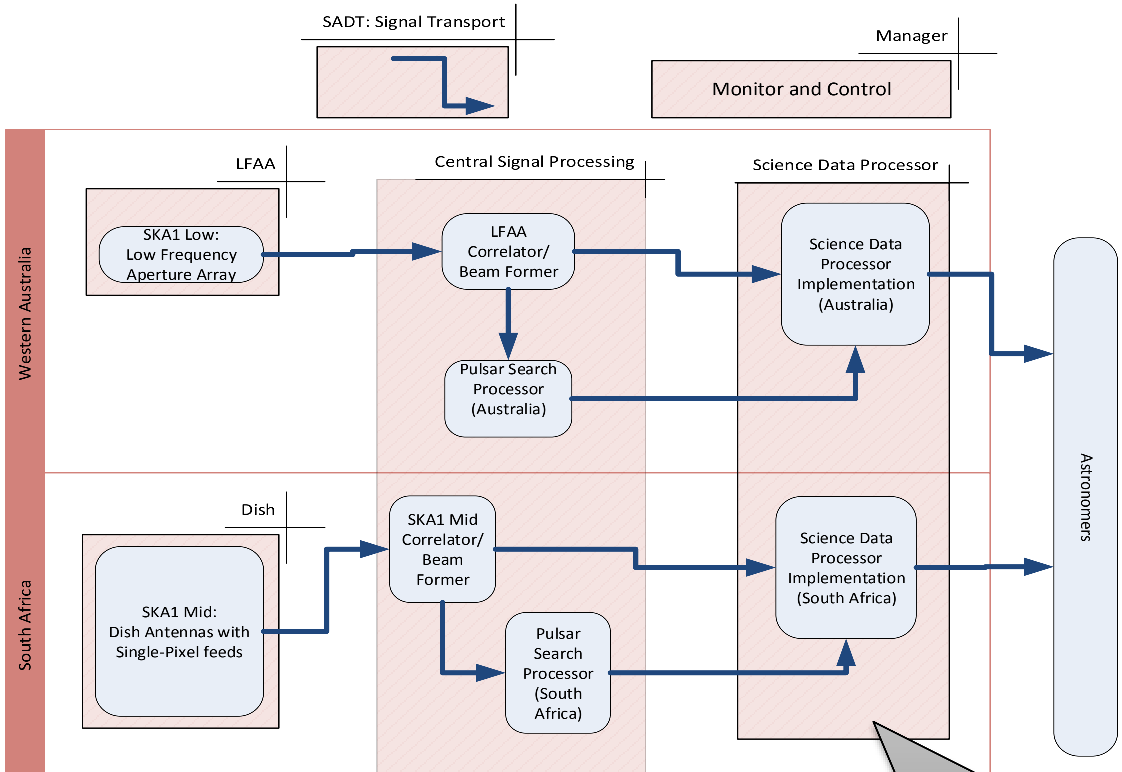

These are off-site! (In Perth & Cape Town)

Role of computing/processing in SKA:

COMPUTING IS THE MAJOR PART OF SKA TELESCOPES BY DESIGN

Large "D" – vs – Large "N"

GBT 100-m diameter telescope

SKA LFAA prototype array

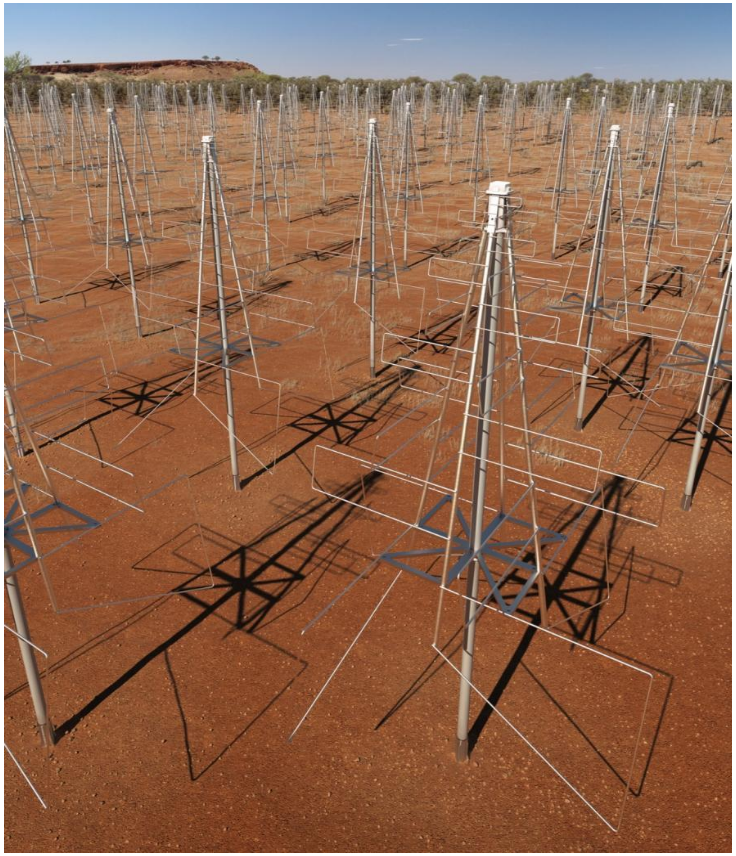

No 1 aim: collect as many photons as possible
No 2 aim: maximum separation of collectors -> achieve high angular resolution

Basics of Interferometry – how signals are combined

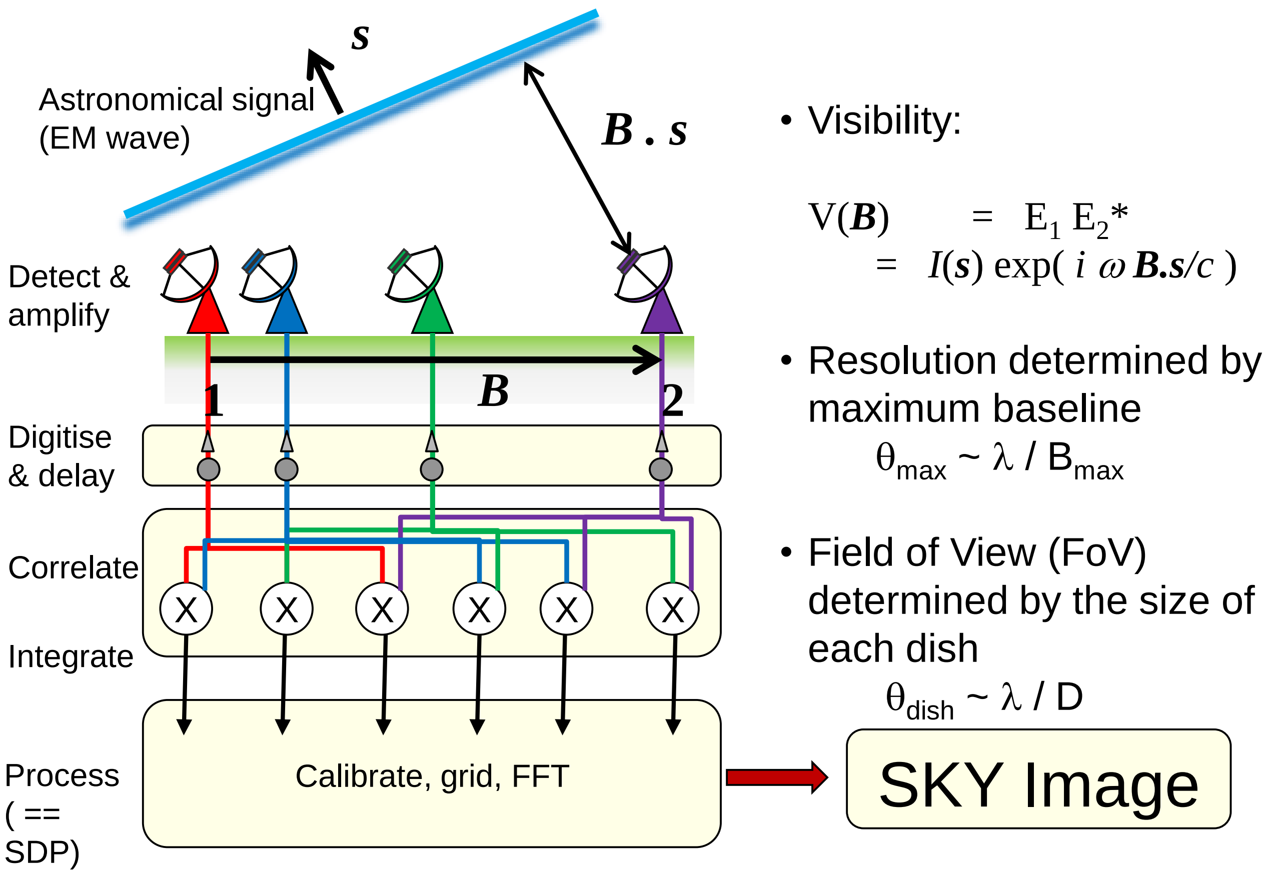

- Visibility:

$$V(\boldsymbol{B}) = E_1 E_2^* = I(\boldsymbol{s}) \exp(i\, \omega\, \boldsymbol{B}.\boldsymbol{s}/c)$$

- Resolution determined by maximum baseline

  $\theta_{max} \sim \lambda / B_{max}$

- Field of View (FoV) determined by the size of each dish

  $\theta_{dish} \sim \lambda / D$

SKA SCIENCE DATA PROCESSOR OVERVIEW

Key Characteristics of SKA Data Processing

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

Noisy Data	Very large data volumes, all data are processed in each observation
Sparse and Incomplete Sampling	Corrected for by deconvolution using iterative algorithms (~10 iterations)
Corrupted Measurements	Corrected by jointly solving for the sky brightness distribution and for the slowly changing corruption effects using iterative algorithms
Multiple dimensions of data parallelism	Loosely coupled tasks, large degree of parallelism is inherently available

SDP Top-level Components & Key Performance Requirements -- **SKA Phase 1**

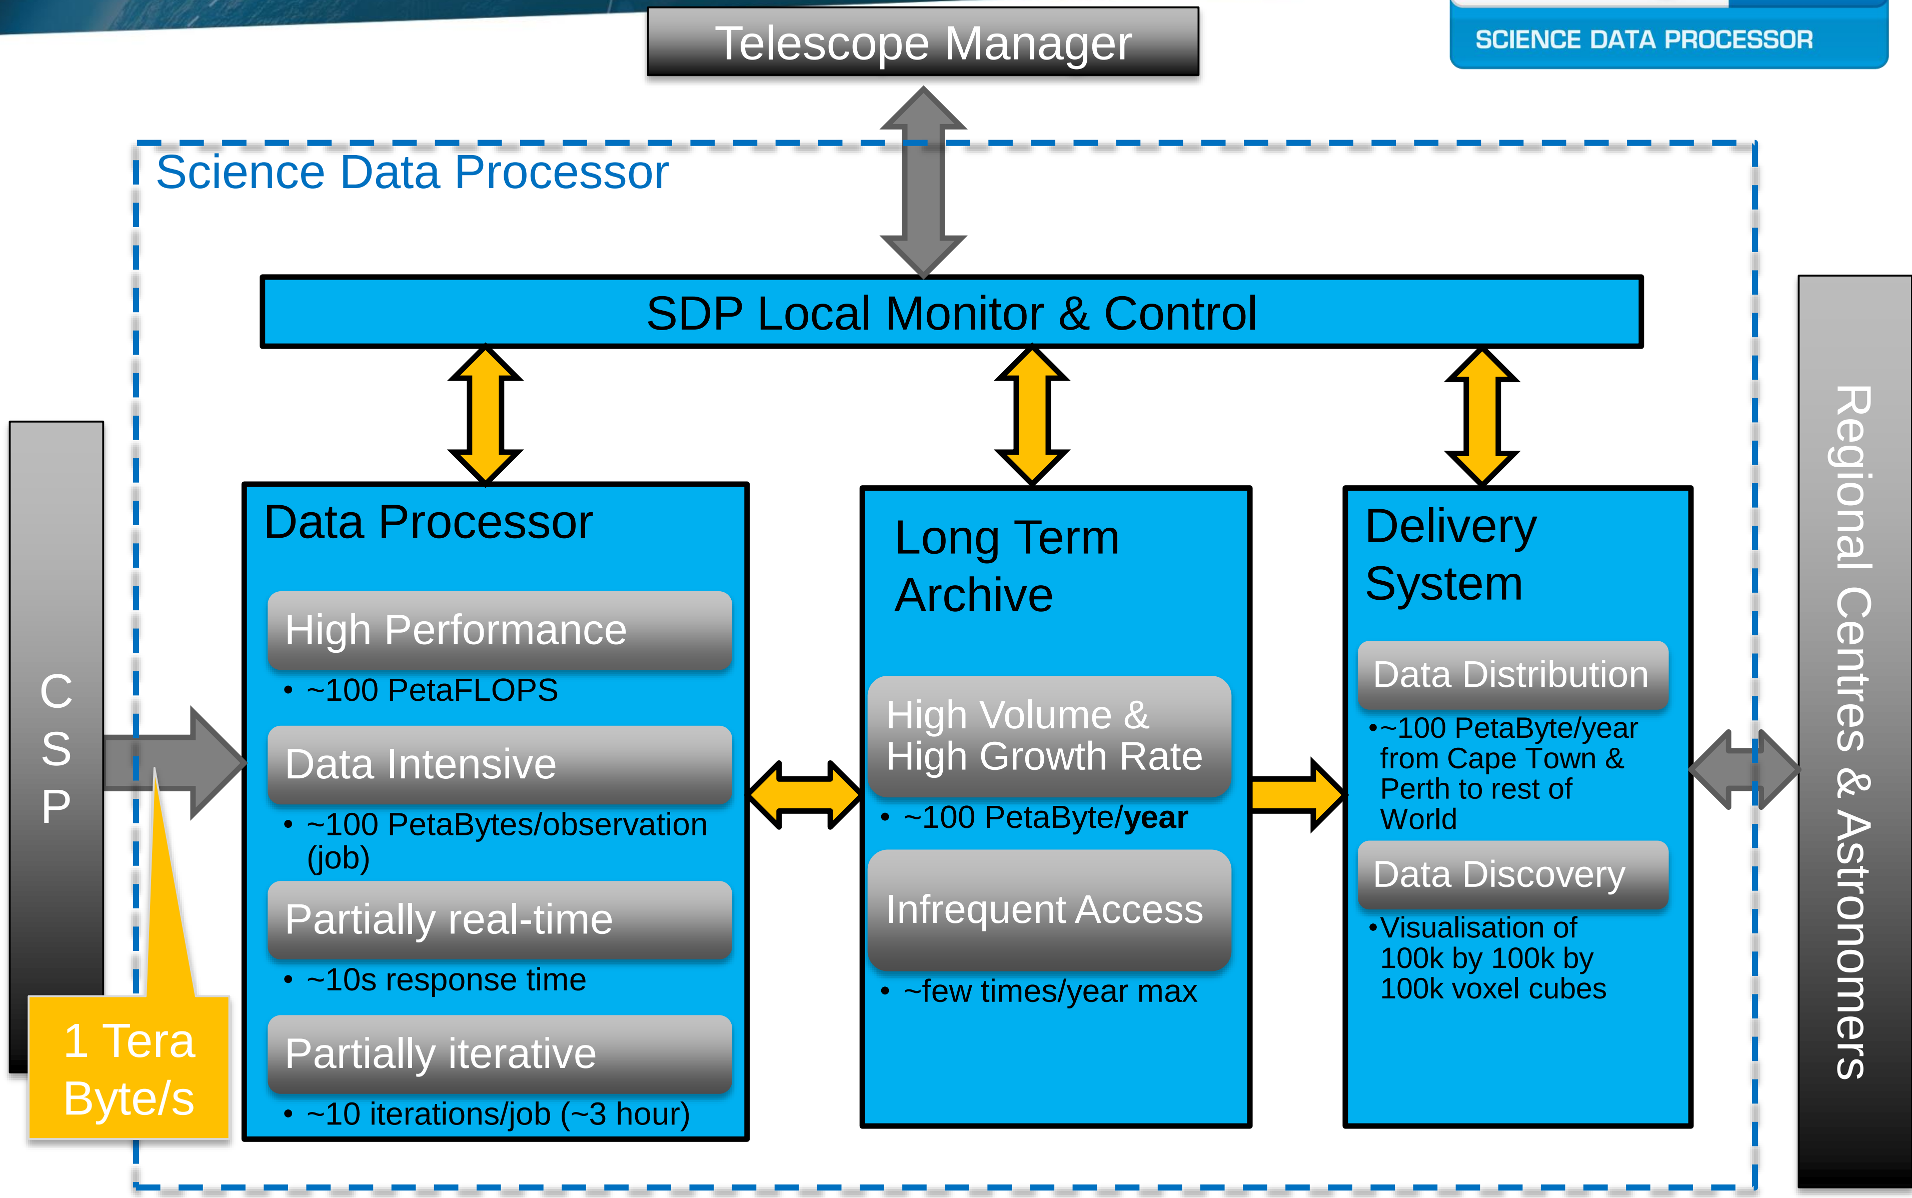

The challenges?

- Power efficiency
  - Funding agencies more tolerant of cap-ex then power op-ex
- Scalability of software
  - Hardware roadmaps indicate h/w will reach requirements
  - Demonstrated software scaling is only 1/1000th of requirement
- Project risks
  - Inevitable significant interaction between software engineers and idiosyncratic domain specific knowledge
  - Software project
- Extensibility, system scalability, maintainability
  - SKA1 is the first “milestone” – expecting significant expansion in the 2020s
  - 50yr observatory lifetime

Relevance of Streaming 2015 technologies

- Multiple Loosely Coupled Task – Yes
  - Within each iteration & between observation
- BSP – Yes: a few (~10) iterations
  - Synchronisation at end of each iteration
  - Relatively small synchronised state (sky model + calibration)
- Workflow – yes
- Streaming – yes (but not interactive)
- Steering – yes (feed back of calibration solutions to the receptors)
- Data Driven – yes & yes
  - In the dataflow execution model sense (i.e., not demand-driven)
  - In the programming model sense (i.e., a declarative description of how to reduce the data are attached to the data themselves)

Why streaming?

- Detecting transient phenomena
  - Need to generate alerts latest about ~10s after receiving the data
- Impractical to store data for later processing
  - Data rate in ~ 1 TeraByte/s
  - -> “stream” but with data chunked by observation

SKA Science Data Processor

ARCHITECTURE HIGHLIGHTS

Programming model

- Hybrid programming model:
  - Dataflow at **coarse-grained** level:
    - About 1 million tasks/s max over the whole processor (-> ~10s – 100s milli second tasks), consuming ~100 MegaByte each
    - **Static scheduling** at coarsest-level (down to “data-island”)
      - Static partitioning of the large-volume input data
    - **Dynamic scheduling** within data island:
      - Failure recovery, dynamic load-balancing
    - *Data driven* (all data will be used)
  - Shared memory model at **fine-grained** level e.g.: threads/OpenMP/SIMT-like
    - ~100s active threads per shared memory space
    - Allows manageable working memory size, computational efficiency

Why?

- Shared memory model essential at fine-grain to control **working memory** requirements
- Dataflow (but all of these are still to be proven in our application):
  - Load-balancing
  - Minimisation of data movement
  - Handling failure
  - Adaptability to different system architectures

Fault Tolerance – “Non-Precious” Data

- Classify arcs in the dataflow graph as precious or *non-precious*
- Precious data are treated in usual way – failover, restart, RAID, etc.
- Non-precious data can be dropped:
  - If they are input to a map-type operation then no output
  - If they are input to a reduction then result is computed without them
- Stragglers outputting non-precious data can be terminated after a relatively short time-out

The non-precious data concept - illustration

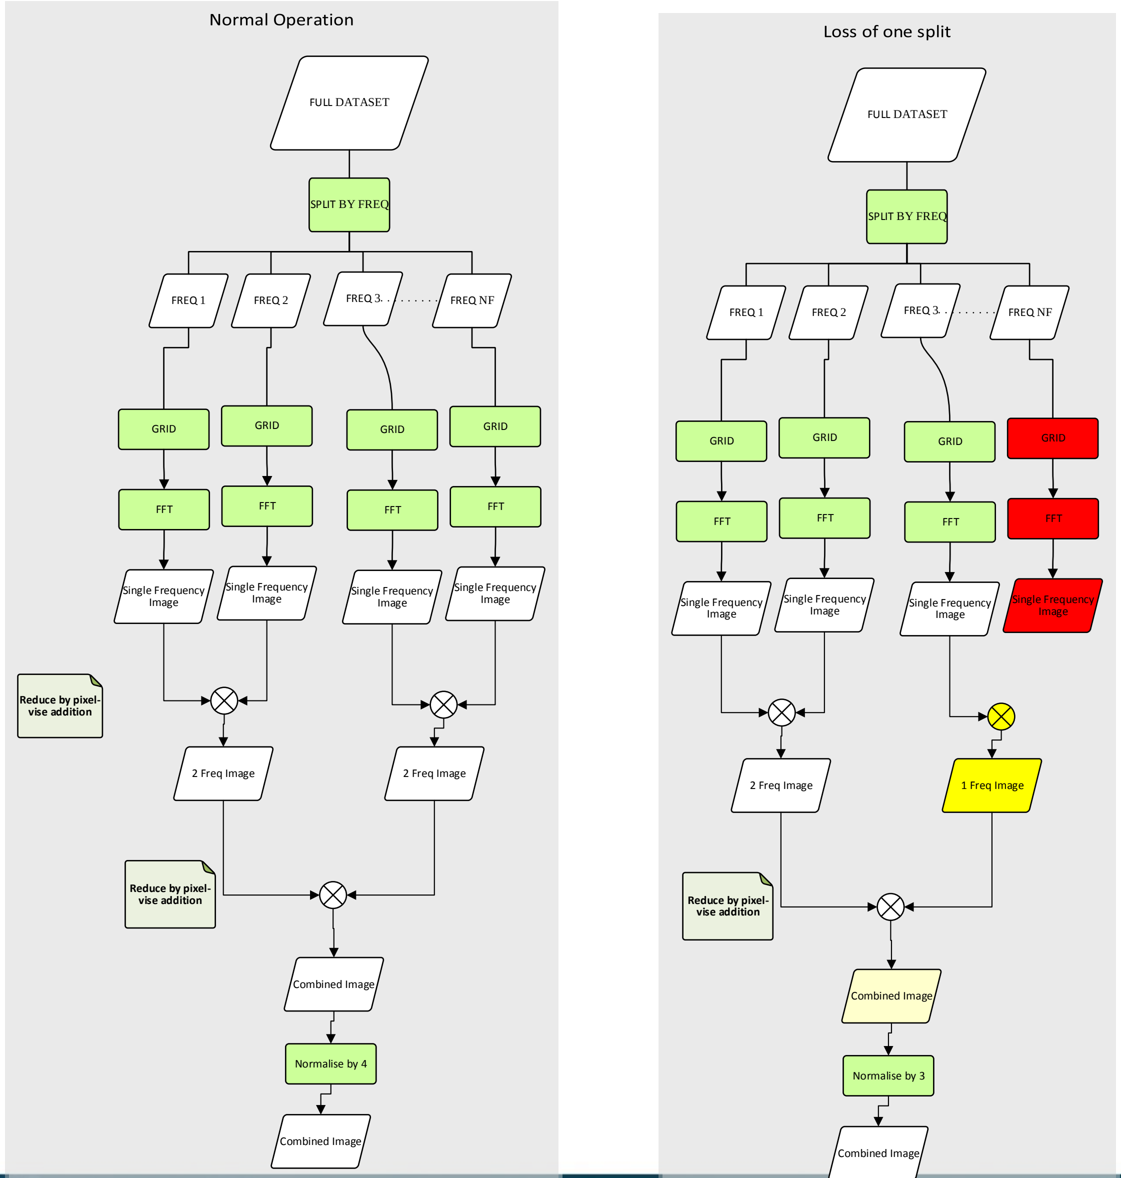

Data Ingress and Interconnect Concept

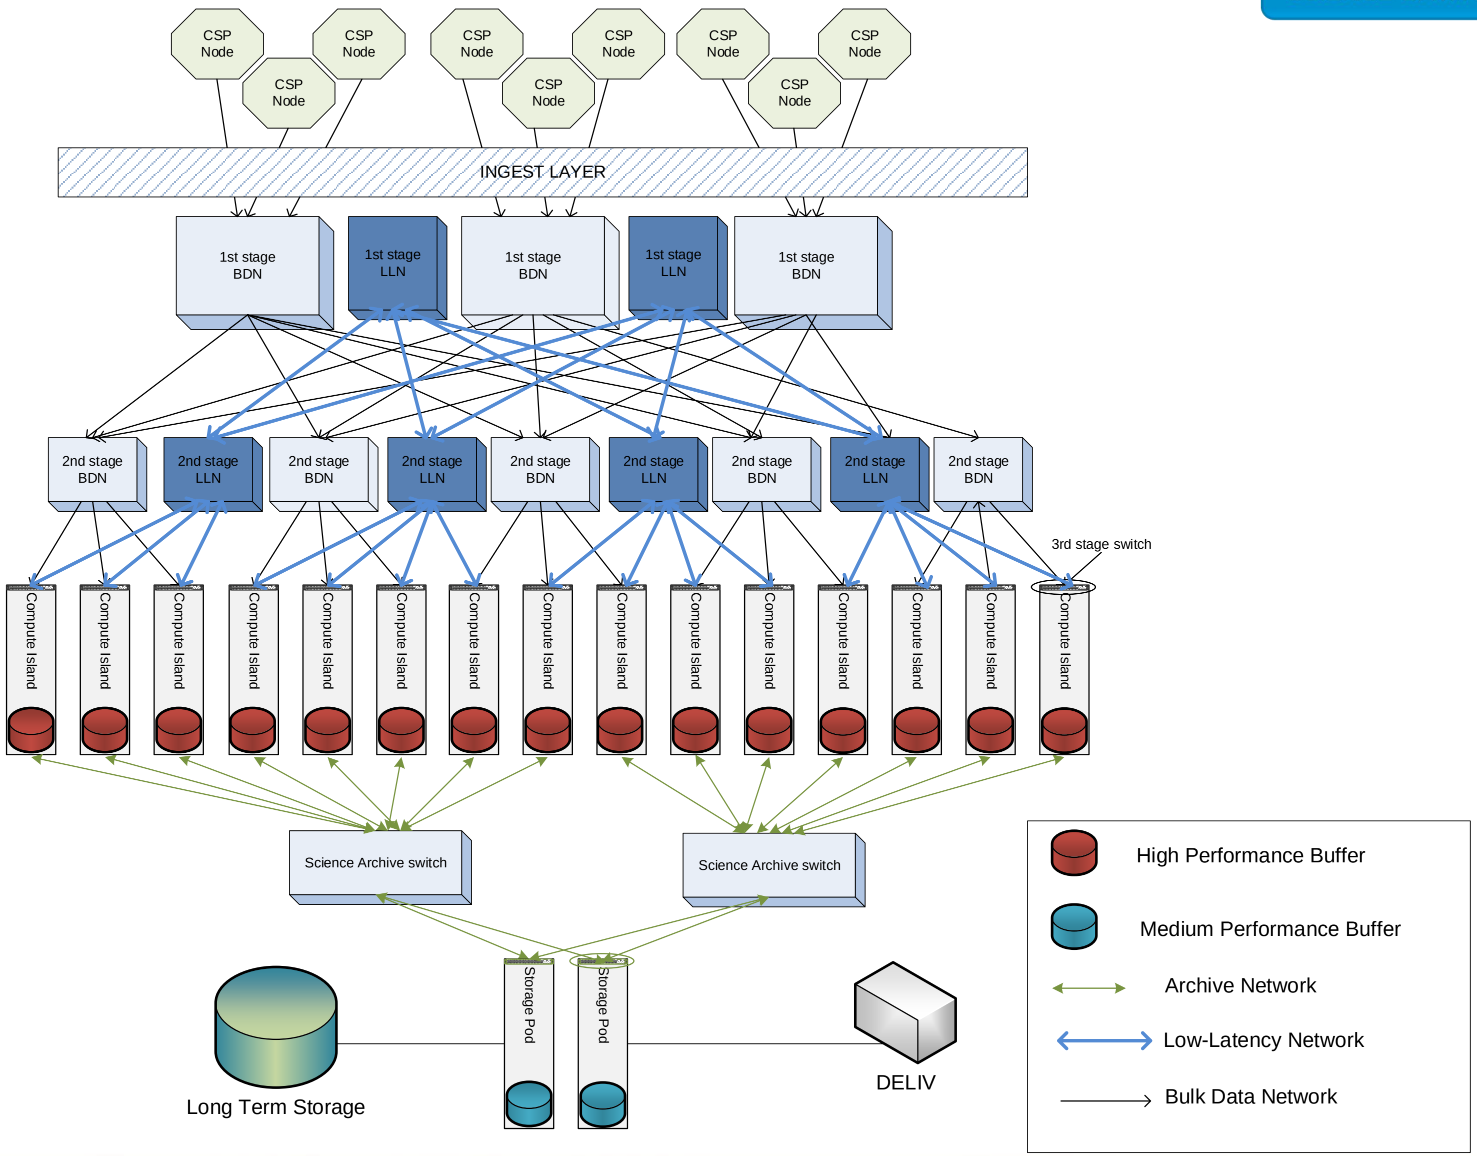

Final Remarks

From SKA-SDP Phase 1 perspective:

- Probably all of the paradigms and ideas we need are available
- Most technologies required for implementing these are available but scattered between many products and platforms
- Need to move towards integrating these together and stable APIs

SKA
SQUARE KILOMETRE ARRAY
SDP
SCIENCE DATA PROCESSOR