The Square Kilometre Array

Networks and Computing

Dr Shaun W Amy | Data Transmission Specialist/Computing Infrastructure Group Leader

6 September 2012

CSIRO ASTRONOMY AND SPACE SCIENCE
www.csiro.au

Overview

- Introduction to Radio Astronomy:
  - some history,
  - important parameters,
- The Australian SKA Pathfinder (ASKAP):
  - overview,
  - site network,
  - long-haul network(s),
  - computing.
- The Square Kilometre Array (SKA):
  - what is it and where will it be located,
  - indicative data and computing requirements.
- Lessons Learnt and Challenges Ahead

A Brief Introduction to Radio Astronomy

The Electromagnetic Spectrum

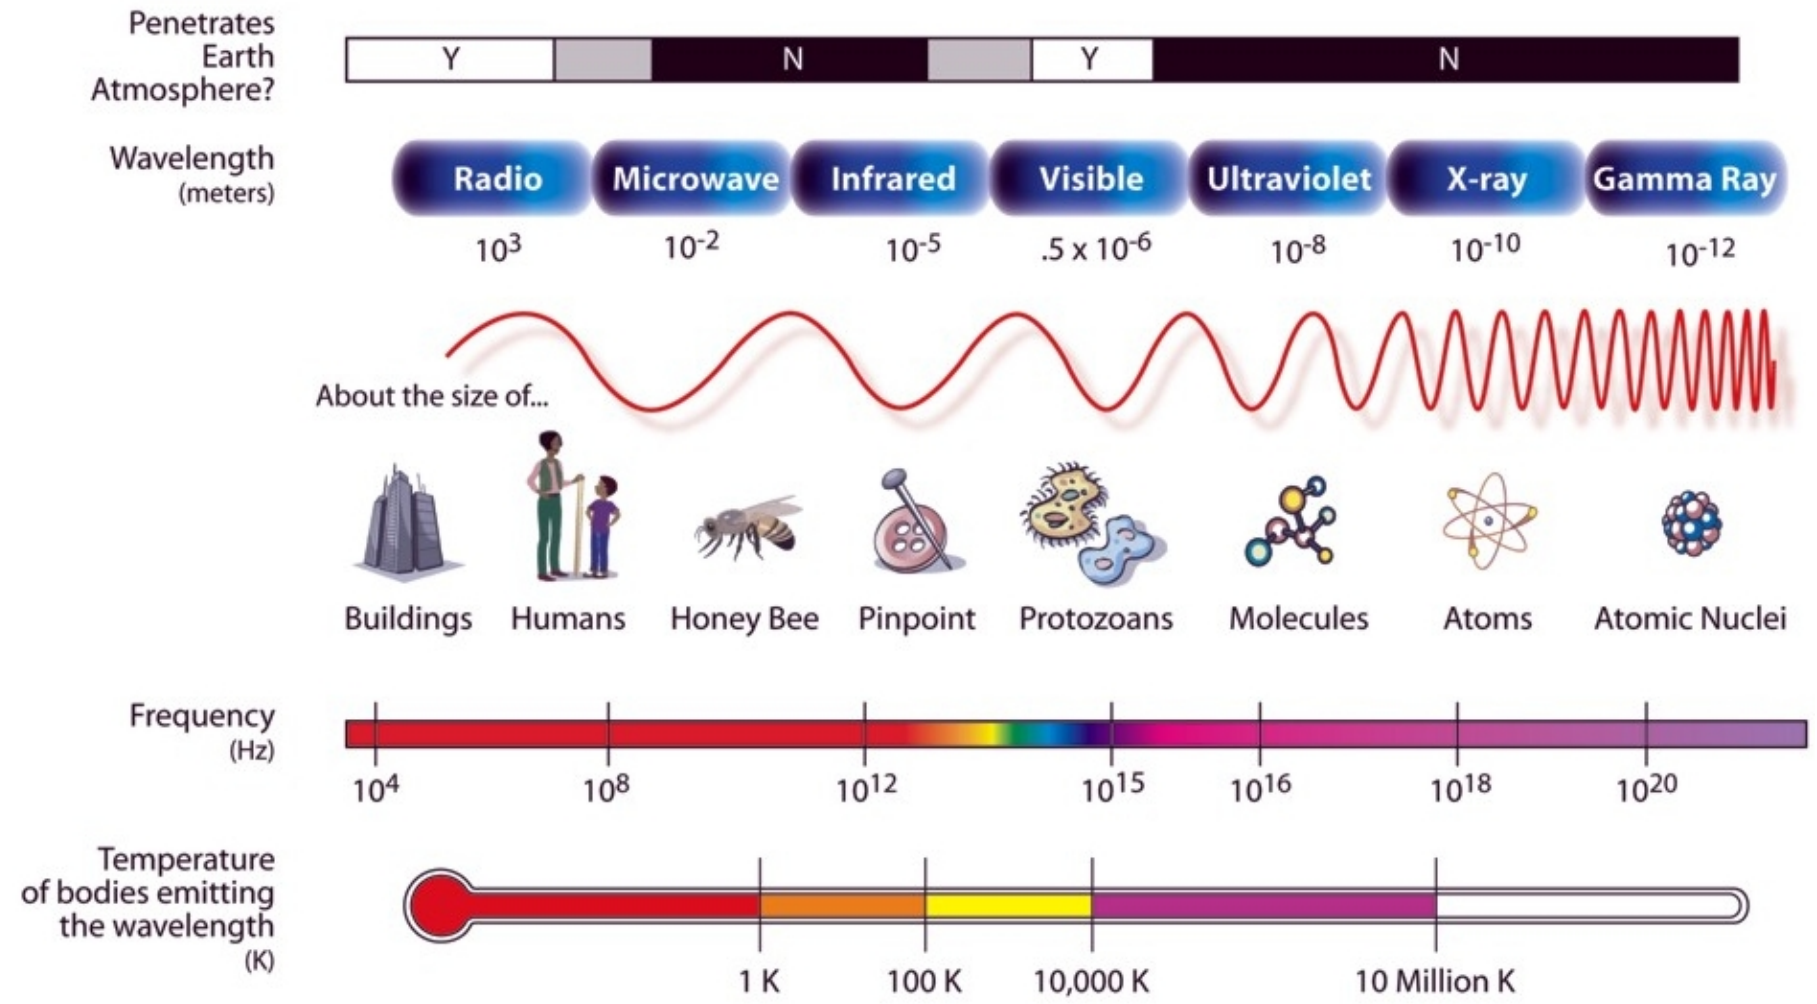

Image: NASA

Cosmic “Noise” (1931 – 1933)

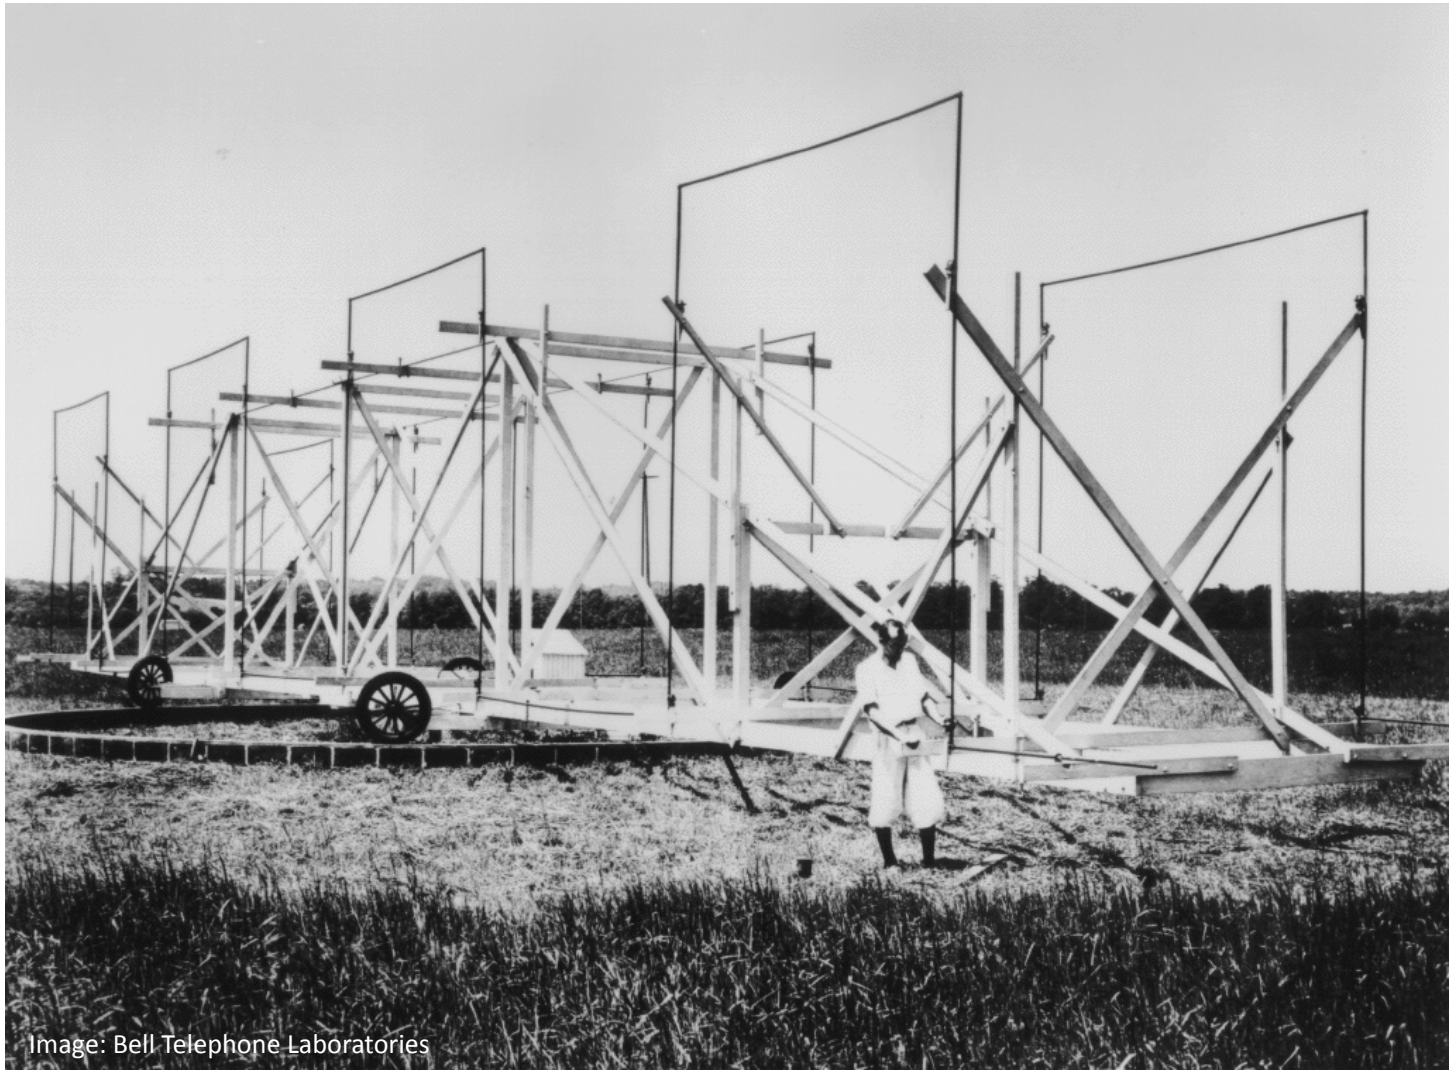

Image: Bell Telephone Laboratories

The Good Ol' Days

- A little over 40 years since the discovery of pulsars by Jocelyn Bell and Tony Hewish at Cambridge.
- Was referred to as "scruff" on the chart recording made at 81.5MHz.
- Deduced it was not from a terrestrial source as the same signal occurred almost 4 minutes earlier each day.
- Confirmed a (radical) prediction made in 1934 by Baade and Zwicky.

Fig. 1.1. Discovery observations of the first pulsar. (*a*) The first recording of PSR 1919+21; the signal resembled the radio interference also seen on this chart: (*b*) Fast chart recording showing individual pulses as downward deflections of the trace.

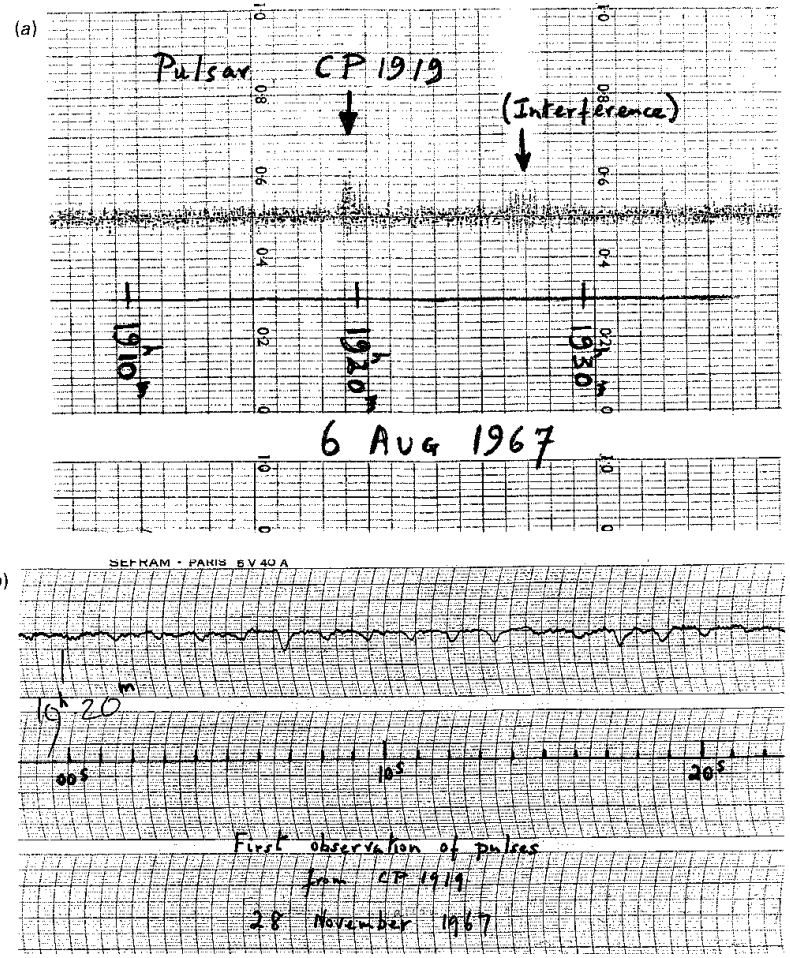

From: Lyne and Graham-Smith: *Pulsar Astronomy* (1990)

Radio Astronomers Want It All…

- Sensitivity
- Angular Resolution
- Spectral Resolution
- Dynamic Range
- Instantaneous Bandwidth
- Frequency Range
- Survey Speed (Field of View)
- Near-infinite computing power and on-line storage
- State of the Art data reduction and analysis tools
- Free of RFI

Resolution = Observing wavelength / Telescope diameter

Angular Resolution	Optical (5000A) Diameter	Optical (5000A) Instrument	Radio (4cm) Diameter	Radio (4cm) Instrument
$1'$	**2mm**	**Eye**	**140m**	**GBT+**
$1''$	**10cm**	**Amateur Telescope**	**8km**	**VLA-B**
$0.''05$	**2m**	**HST**	**160km**	**MERLIN**
$0.''001$	**100m**	**Interferometer**	**8200km**	**VLBI**

Atmosphere gives 1" limit without corrections which are easiest in radio

Jupiter and Io as seen from Earth

1 arcmin | 1 arcsec | 0.05 arcsec | 0.001 arcsec

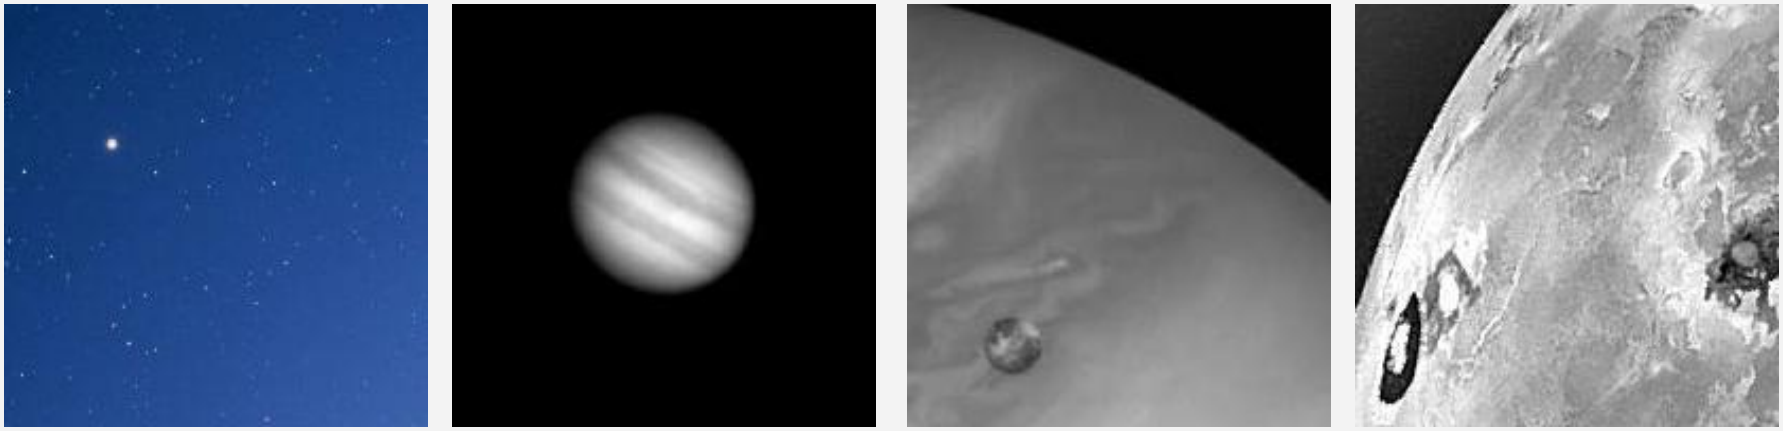

Simulated with Galileo photo

Image: Shaun Amy

Aperture Synthesis Imaging

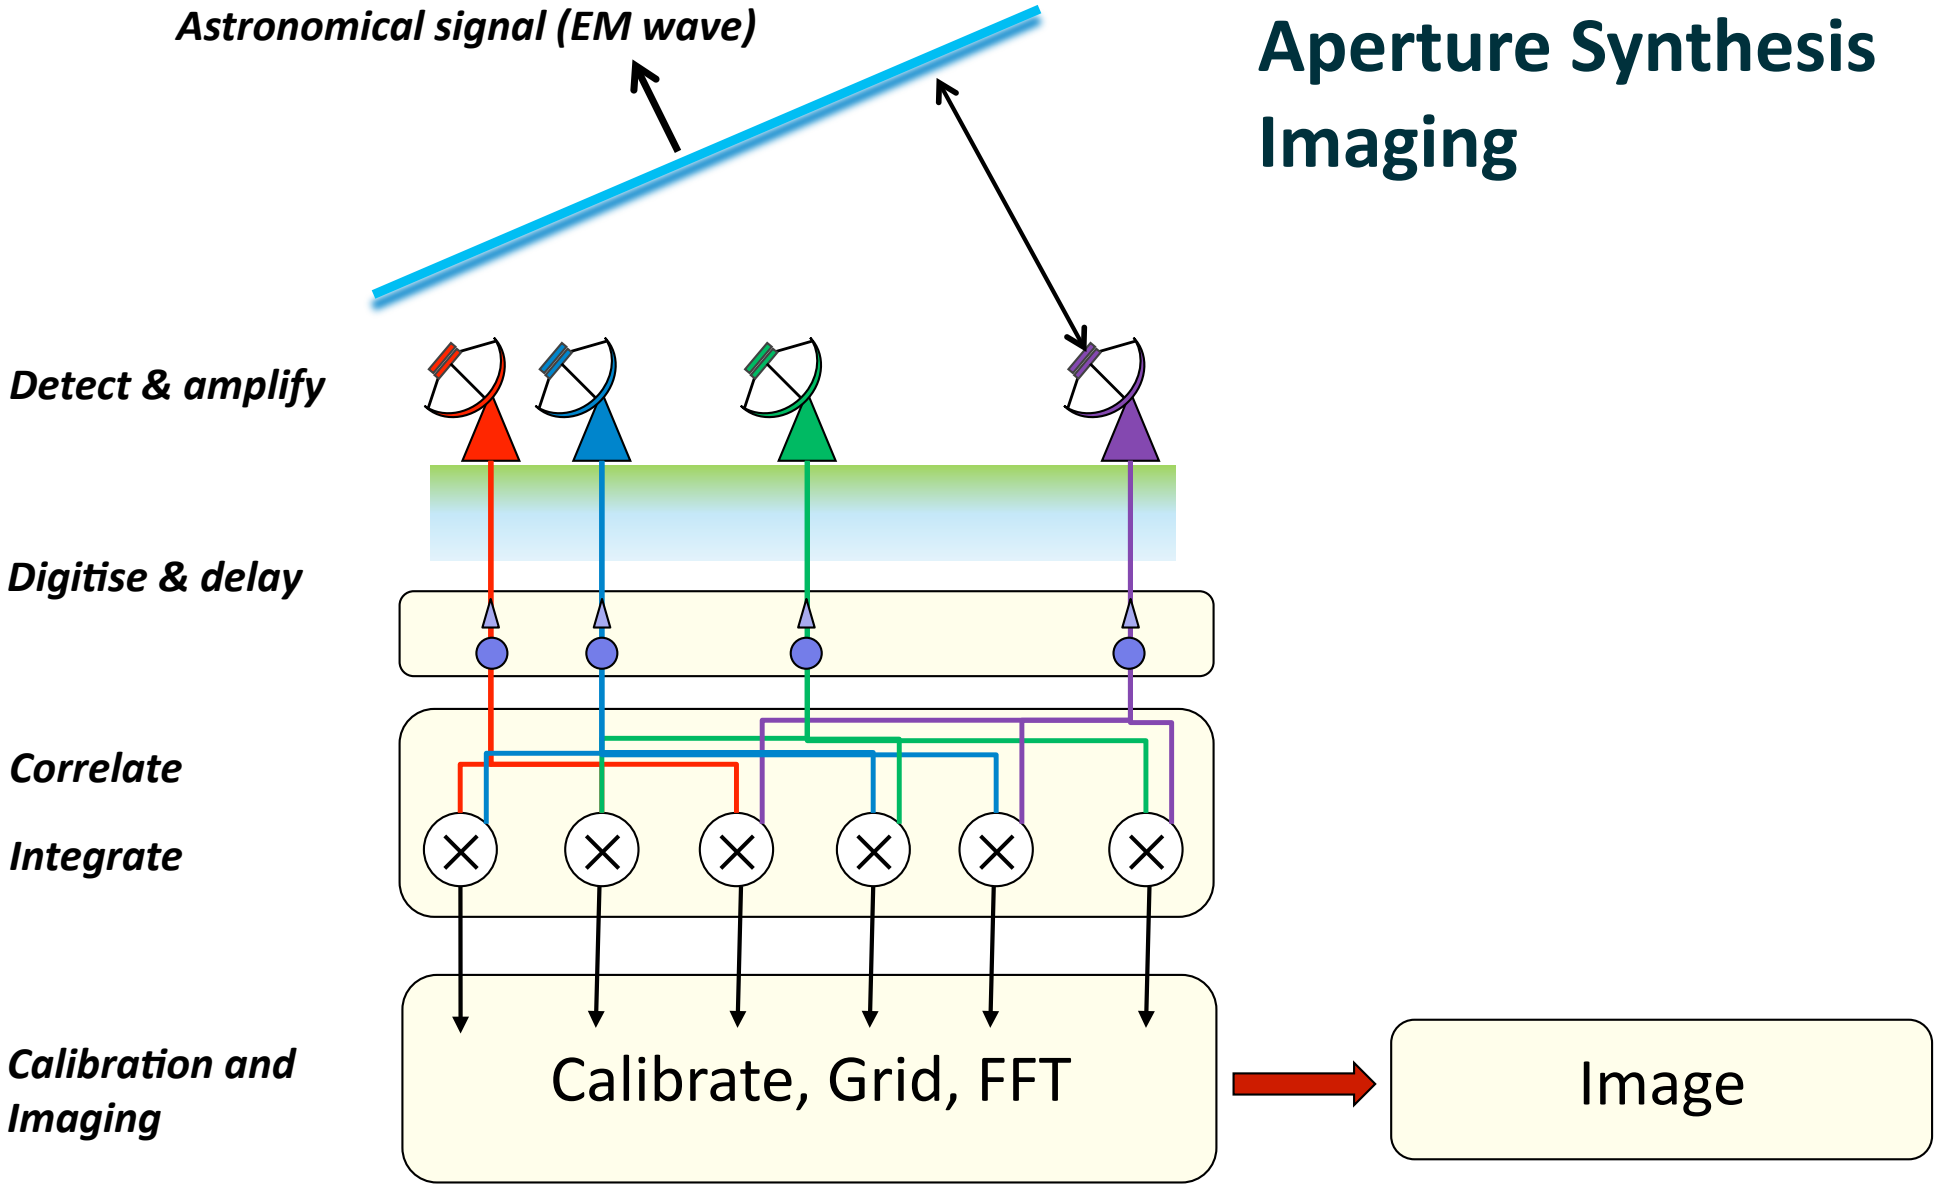

Image: Paul Alexander, University of Cambridge

Image: Shaun Amy


Image: NAIC

Australian SKA Pathfinder (ASKAP)

Image: CSIRO

Australian SKA Pathfinder (ASKAP)

- Sited at the Murchison Radio Observatory, Western Australia.
- Frequency: 0.7 – 1.8 GHz.
- 36 antennas, 12m diameter.
- Started construction: July 2006.
- Official opening: 5 October 2012.

- Data rate from correlator ~2.5GB/s
  - a DVD every two seconds!

- Science processing requirement:
  - 100TF/s for basic capabilities,
  - 400+TF/s for high angular resolution spectral line imaging.


Image: Shaun Amy

Phased Array Feeds

- Key technology development:
  - increases survey speed by approximately an order of magnitude.
- 188 Receiver Elements:
  - typical radio-telescope has two (single-pixel feed, dual polarisation),
  - data-rate ~1.9Tbit/s from each antenna.

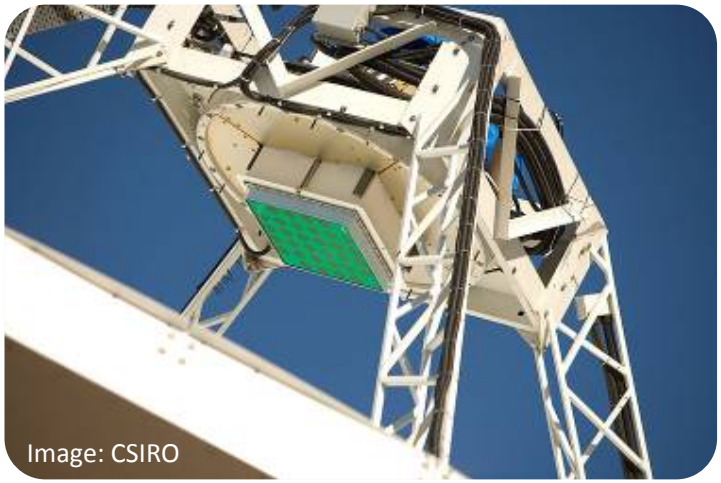
Image: CSIRO

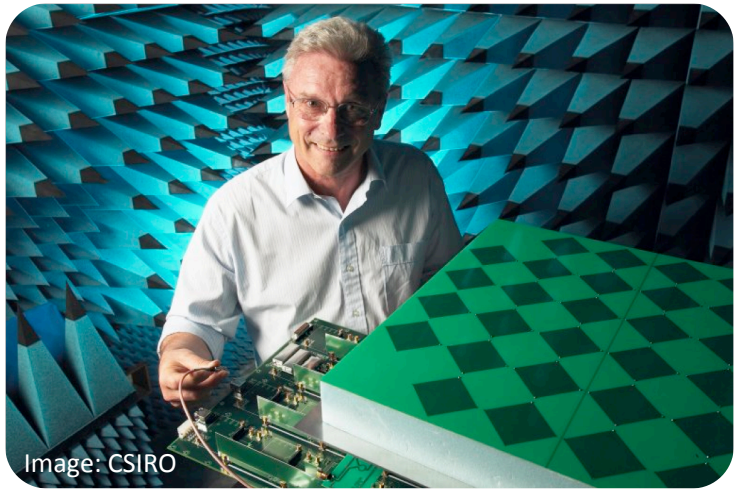
Image: CSIRO

PAF Synthesis Imaging

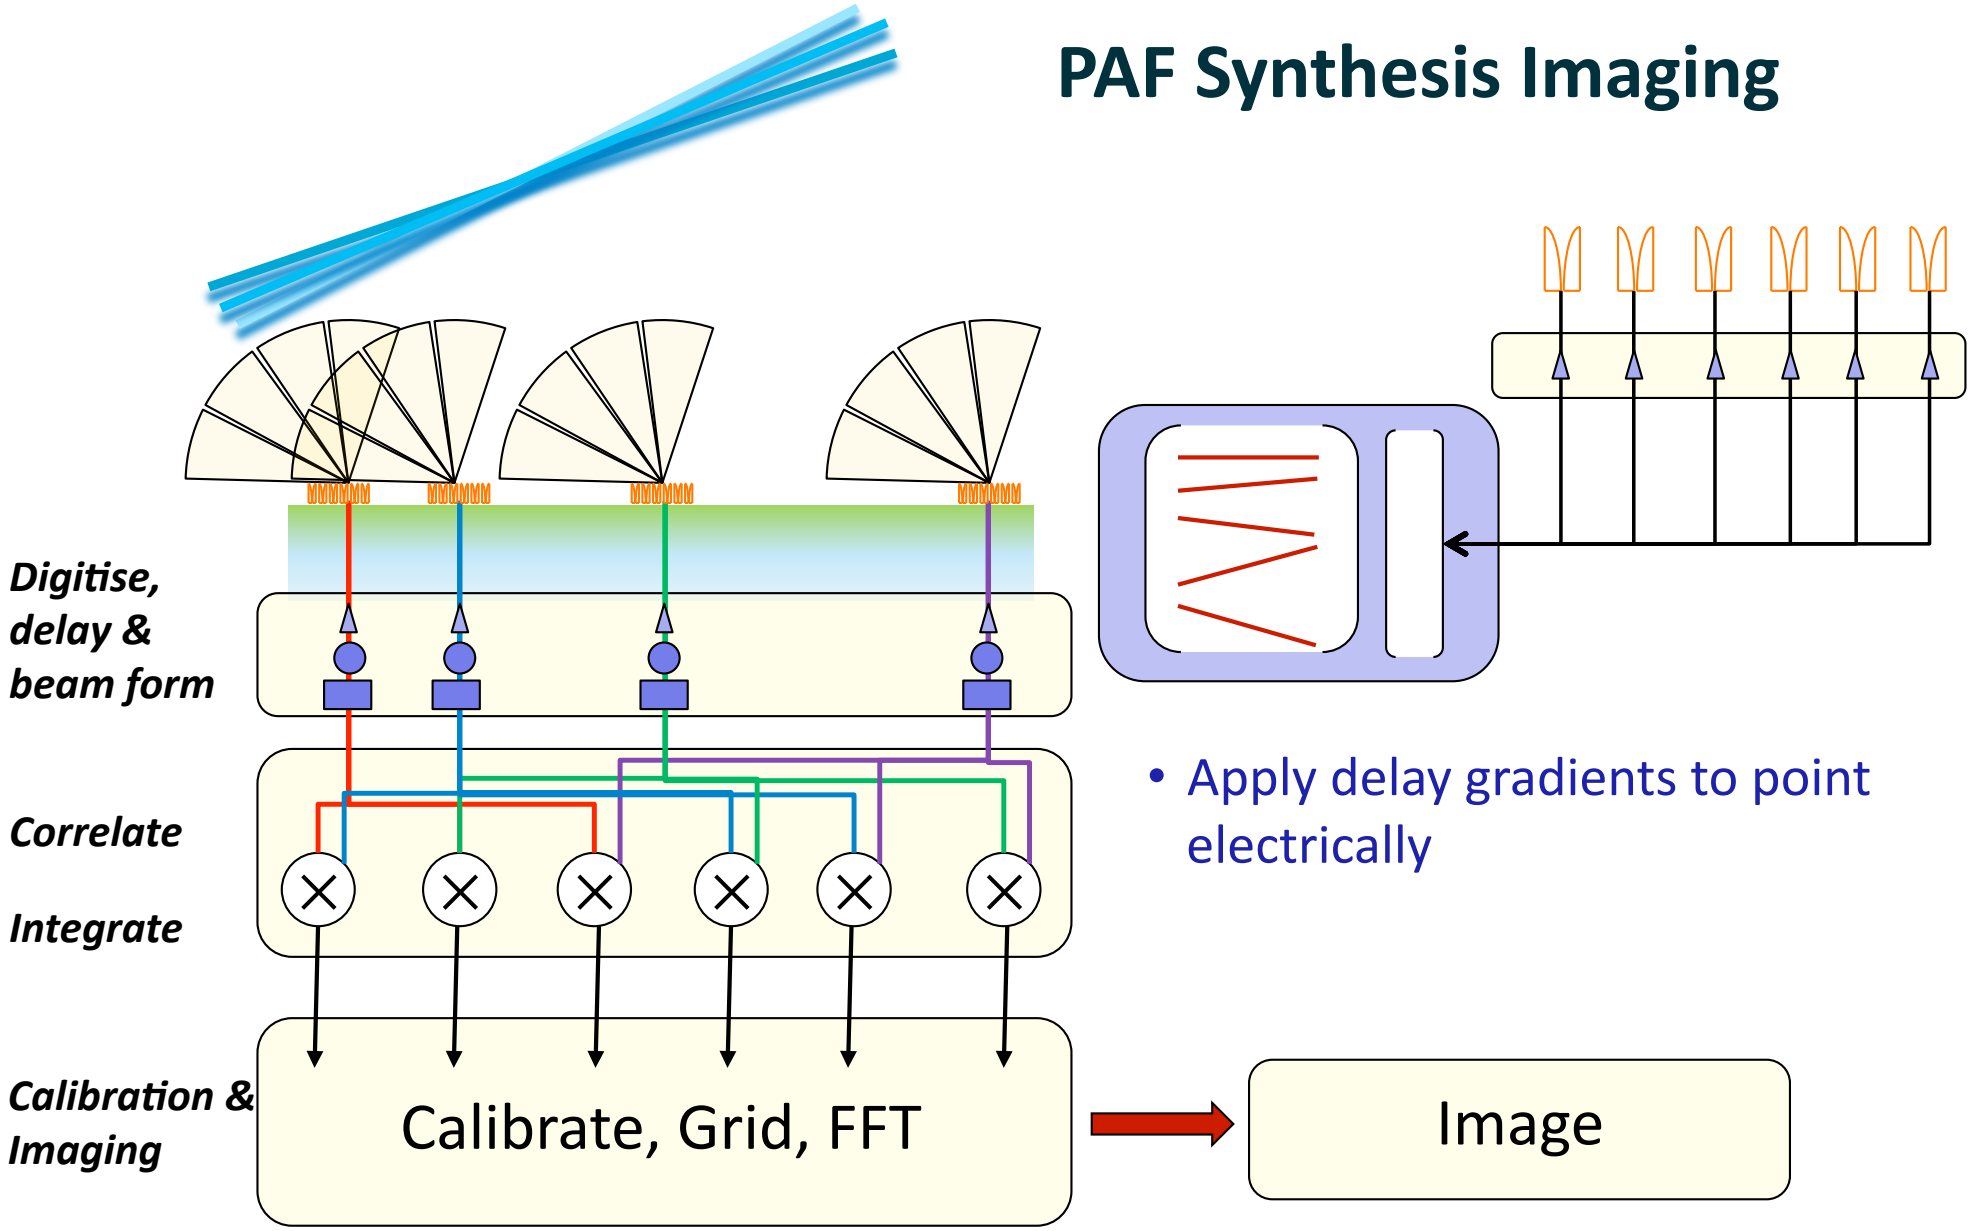

- Apply delay gradients to point electrically

Image: Paul Alexander, University of Cambridge

Digitised RF Data Transmission (per antenna)

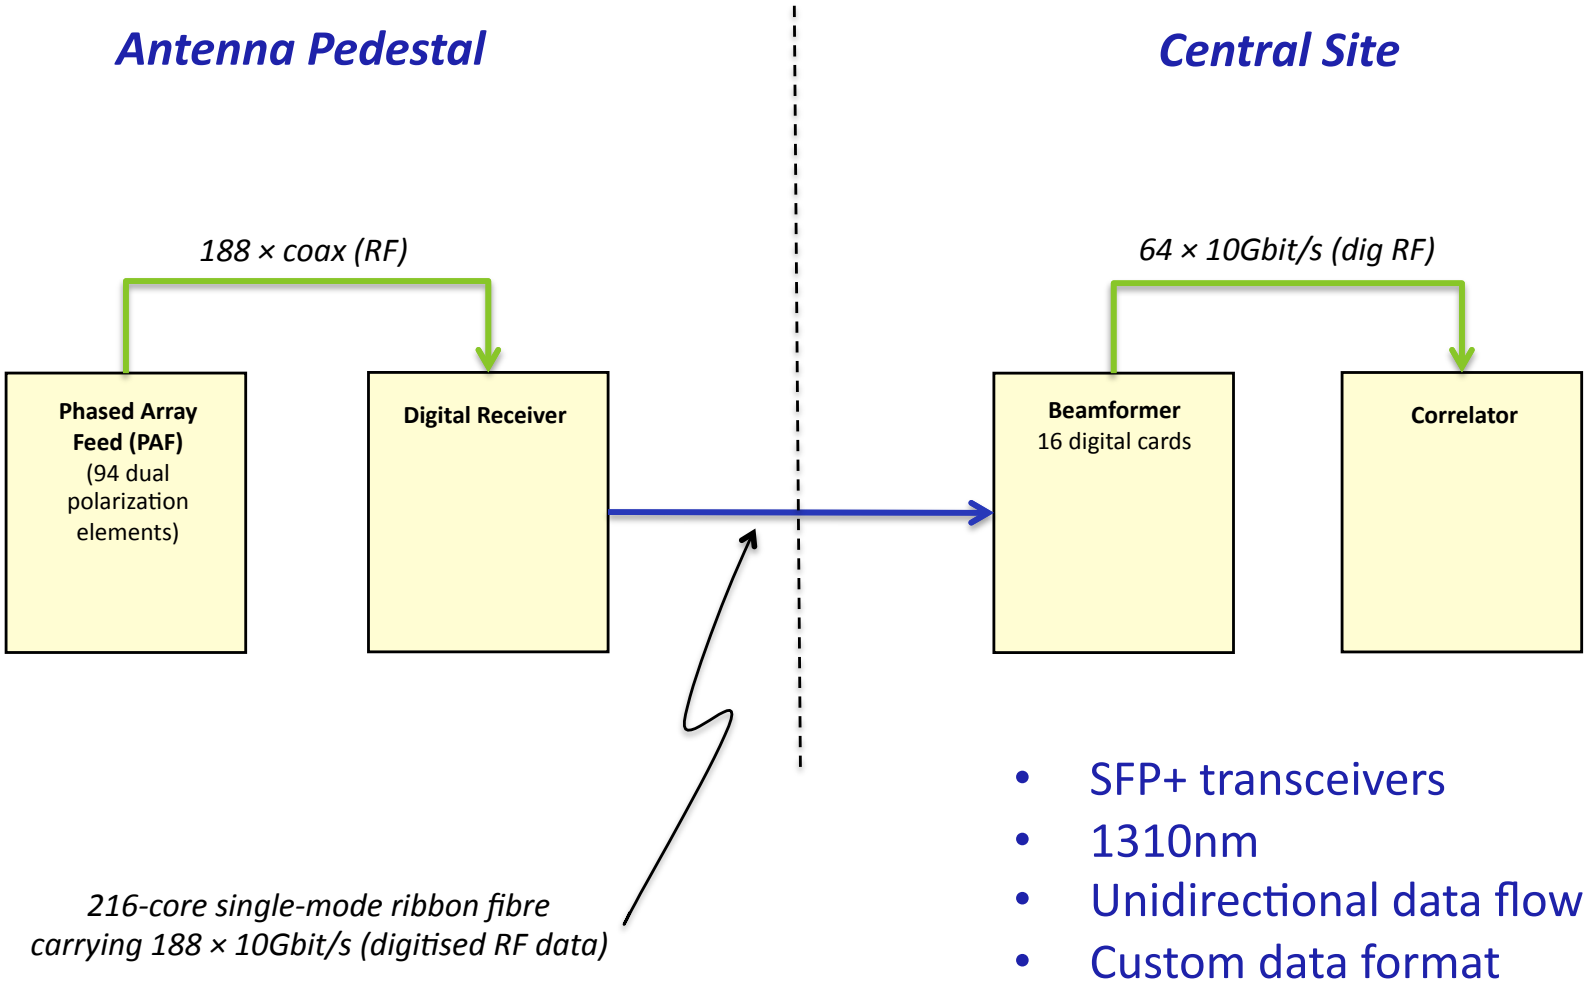

- SFP+ transceivers
- 1310nm
- Unidirectional data flow
- Custom data format

MRO Site Network


Image: CSIRO

- 216 core “ribbon” fibre from each antenna to the control building


Image: CSIRO

MRO Site Network Considerations

- The network core to handle the telescope data will use a tiered data centre class switch approach:
  - Cisco NEXUS 7010,
  - Cisco NEXUS 5596UP + Fabric Extenders (FEX),
  - storage array and hosts use fibre channel.
- Need to consider carefully over-subscription:
  - most equipment has some over-subscription,
  - need to understand the data flows (sustained and peak) to ensure no data loss (mainly UDP):
    - each 1Gbit/s port can produce about 800Mbit/s of data,
    - data flows are almost all uni-directional.
- Scaled-down system for BETA (first six antennas):
  - NEXUS 5596UP + FEX,
  - small number of 1RU switches for environmental monitoring etc.

The ASKAP BETA Network

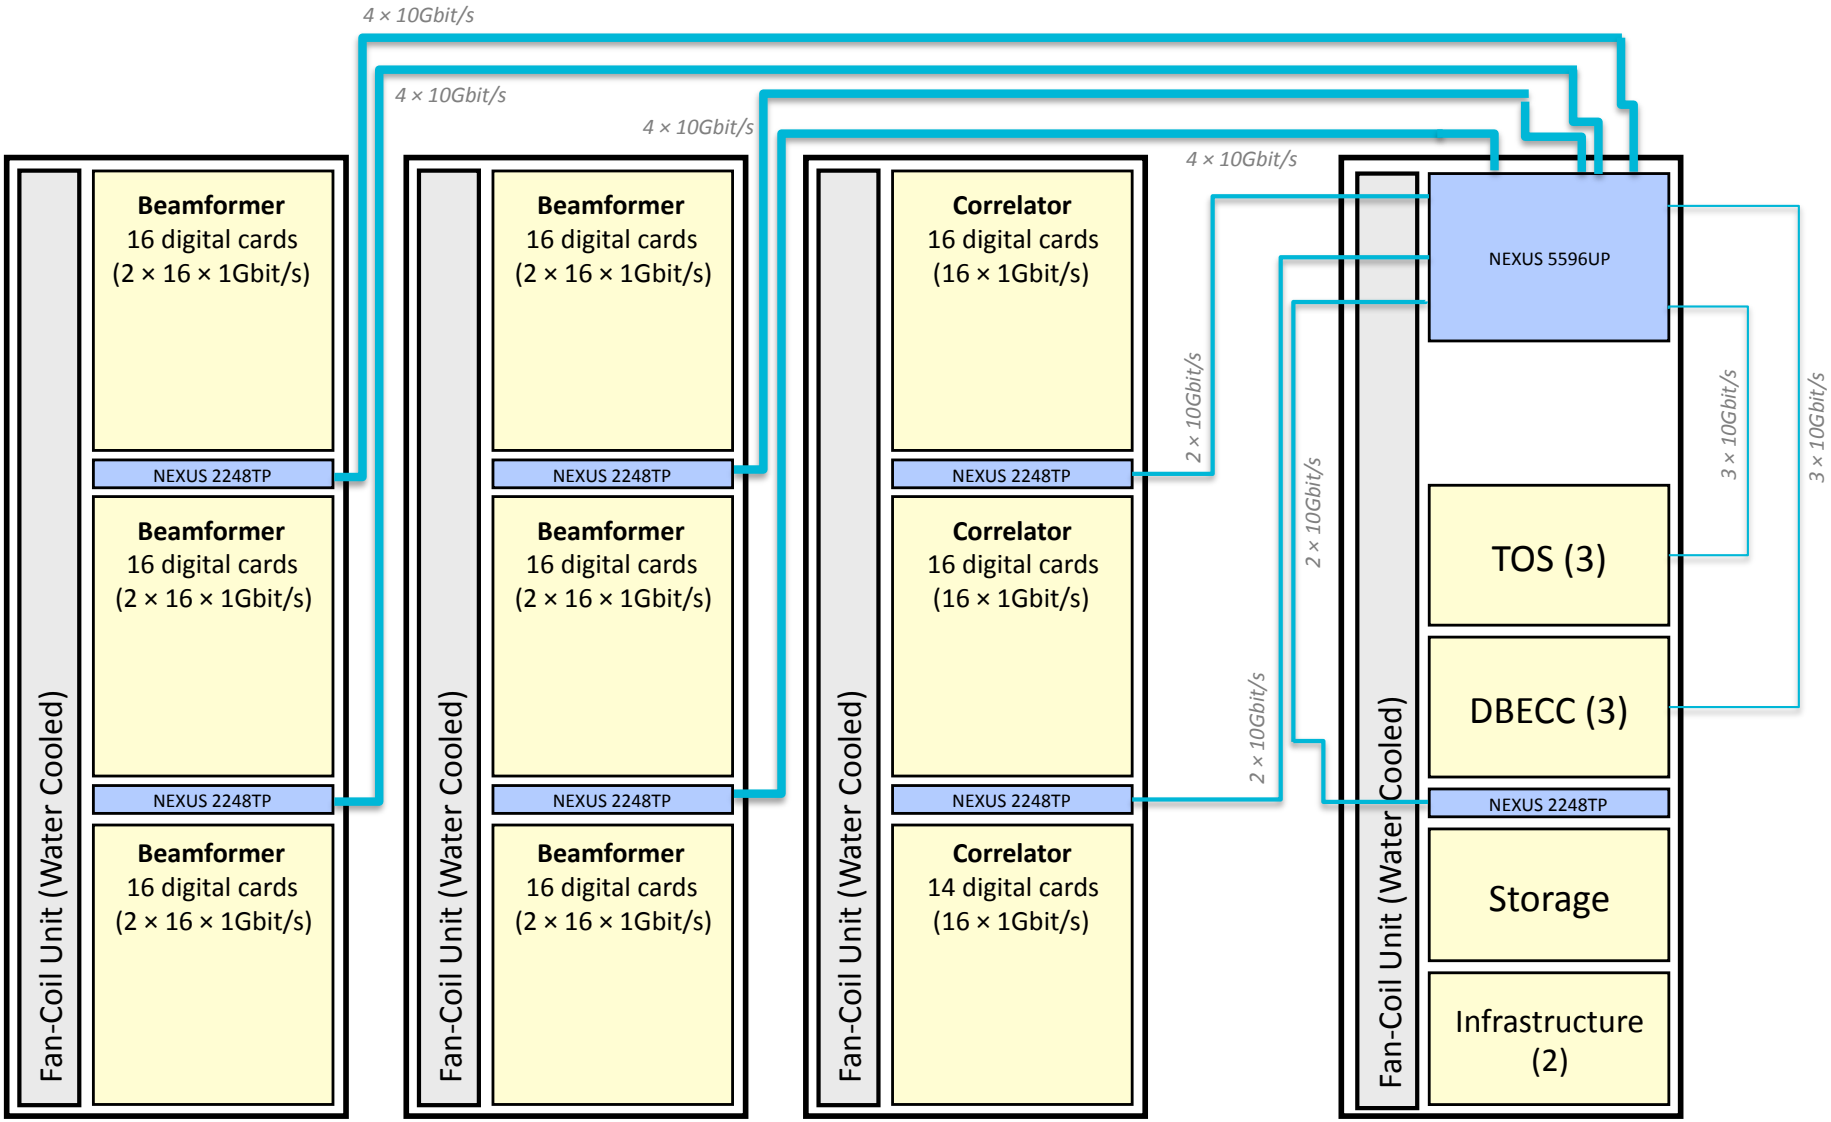

Note: Beamformer – Correlator Data is via a direct non-Ethernet connection between ATCA chassis

MRO Long-Haul Network: Fibre

- Geraldton – MRO:
  - SKA ready,
  - significant construction project (telco standard),
  - 48-core G.652 ULL (72 cores between Geraldton and Mullewa),
  - three “repeater” sites:
    - Mullewa (grid power),
    - Yuin Station (solar power, passively cooled),
    - Murgoo Station (solar power, passively cooled).
- Geraldton – Perth:
  - RBBS project, now part of the NBN,
  - additional fibre installed, nominally for “SKA use”,
  - access agreements complex because many parties involved.

A Solar Powered, Passively Cooled Shelter

Image: Shaun Amy

Murgoo CEV Environmental Performance

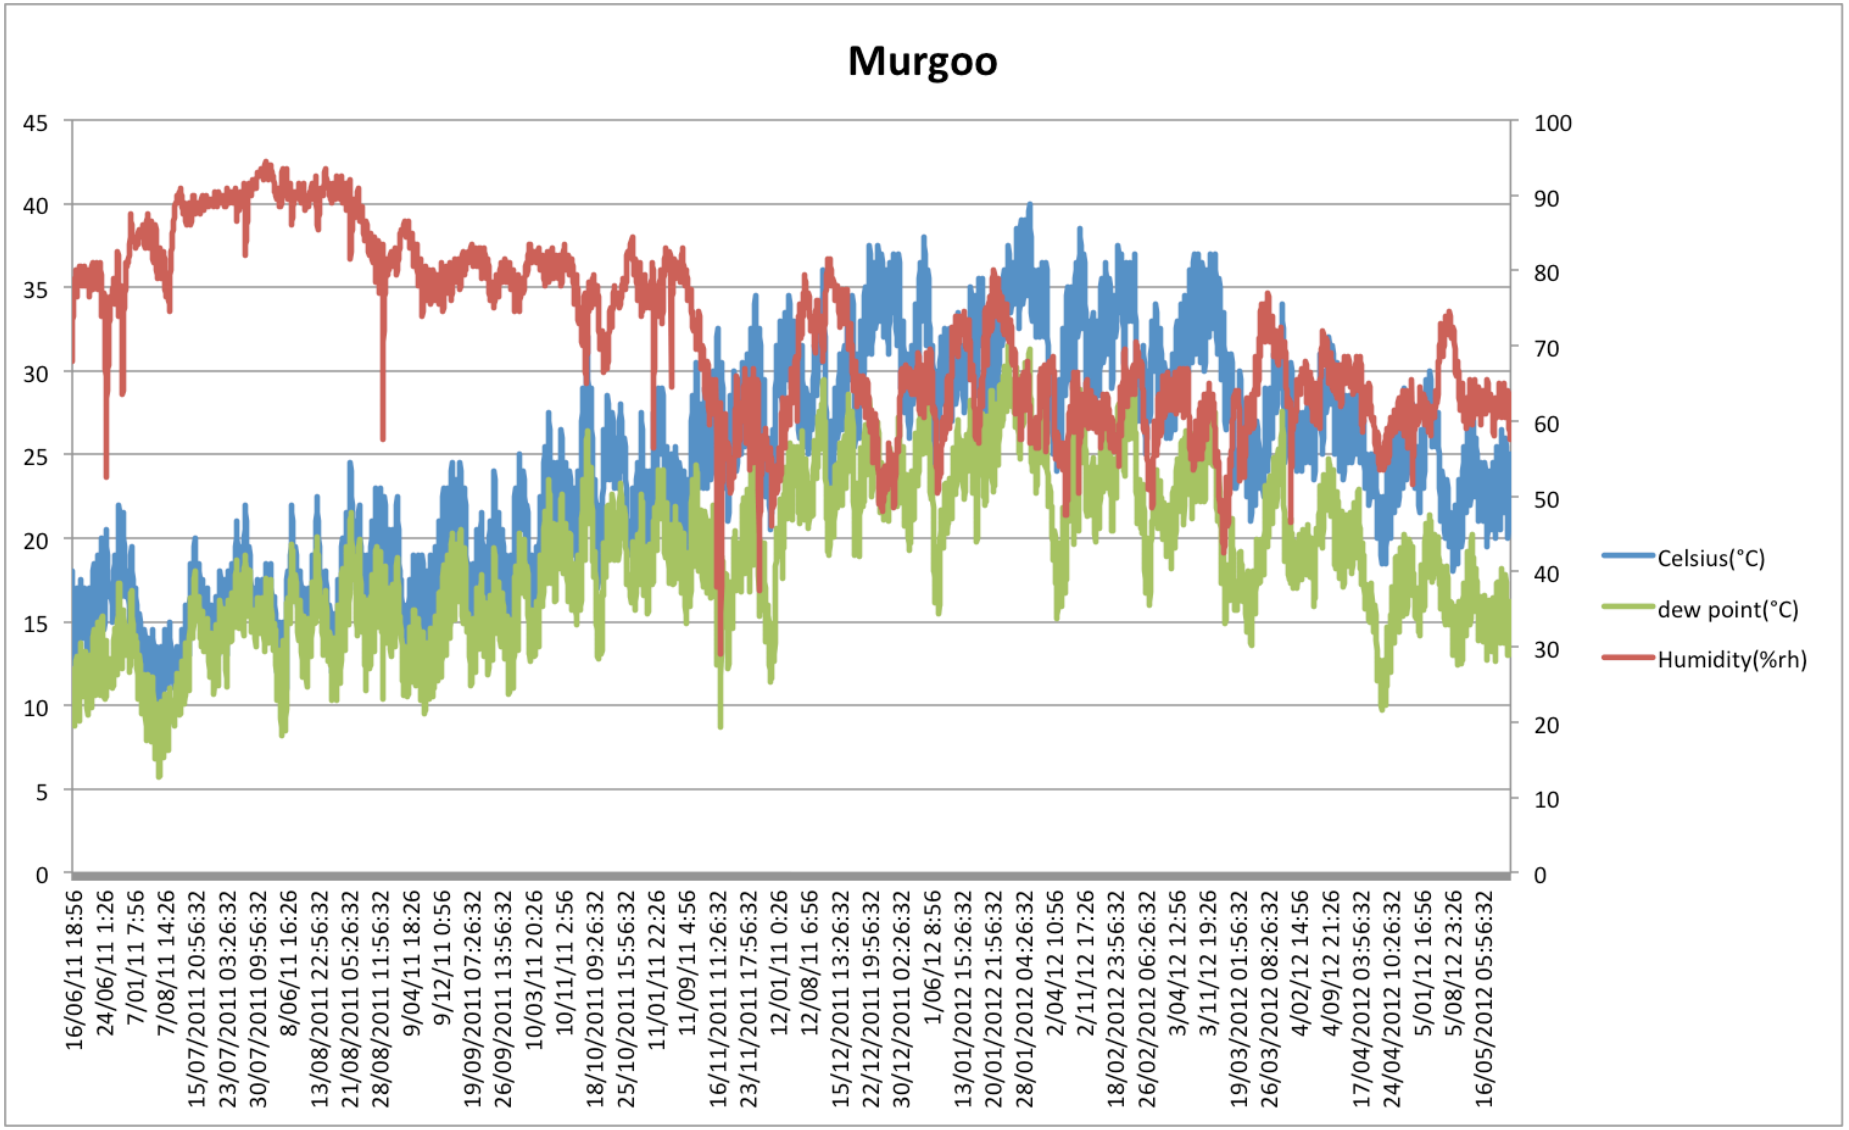

MRO Long-Haul Network: Active Network

- Two networks:
  - 2 × 1Gbit/s “Christmas-Tree Lights” network:
    - simple to get operational,
    - provision of “en-route” services,
    - enabled important initial demonstrators and science to be undertaken,
    - end-to-end confirmation of network operation.
  - DWDM high-bandwidth network:
    - joint CSIRO/AARNet design,
    - optical amplification: Perth – MRO,
    - add/drop in Geraldton,
    - extensible design, initially 40Gbit/s per channel, easily upgradable to 100Gbit/s (80/96 channels),
    - client-side connectivity is 10Gbit/s Ethernet,
    - bandwidth for co-located instrumentation.

The Pawsey High Performance Computing Centre for SKA Science

- AUD$80M super-computing centre
- 1.2 Petaflops processing capability
- 50 PB near-line (HSM) storage
- Equipment installation begins March 2013

The Pawsey High Performance Computing Centre for SKA Science

- Cray Cascade Supercomputer:
  - ~4600 × Intel Xeon (Ivy Bridge) CPUs,
  - ~200 × Intel Xeon Phi Accelerators,
  - Aries Interconnect,
  - ~7PB Lustre Filesystem (Cray Sonexion).

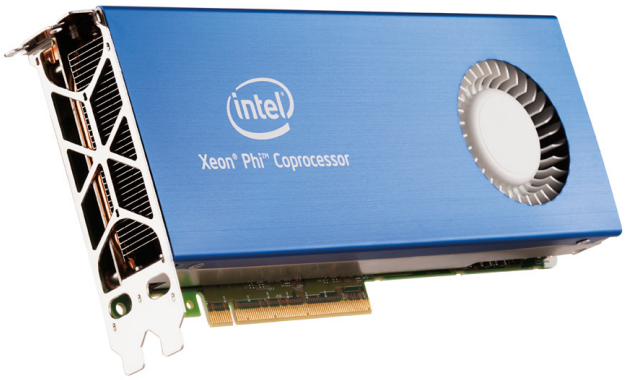

- ~50PB Oracle (Storagetek) Tape Libraries
- SGI Hierarchical Storage Management (HSM) filesystem
- 40Gbit/s Bandwidth to observatory (for ASKAP) + 10Gbit/s (for MWA)
- 10Gbit/s bandwidth (40Gbit/s future) to AARNet:
  - firewall,
  - VPN.

The Square Kilometre Array

Square Kilometre Array

- The SKA will be a revolutionary radio telescope made of **thousands of radio receptors**, or antennas, linked together **across an area the size of a continent.**

- The total collecting area of all the SKA antennas combined will be approximately **one square kilometre.**

- Target cost €1.5 billion.

Members of the SKA Organisation

- Australia
- Canada
- China
- Italy
- New Zealand
- Expect more to join…
- Republic of South Africa
- Sweden
- The Netherlands
- United Kingdom
- India (Associate member)

SKA Core

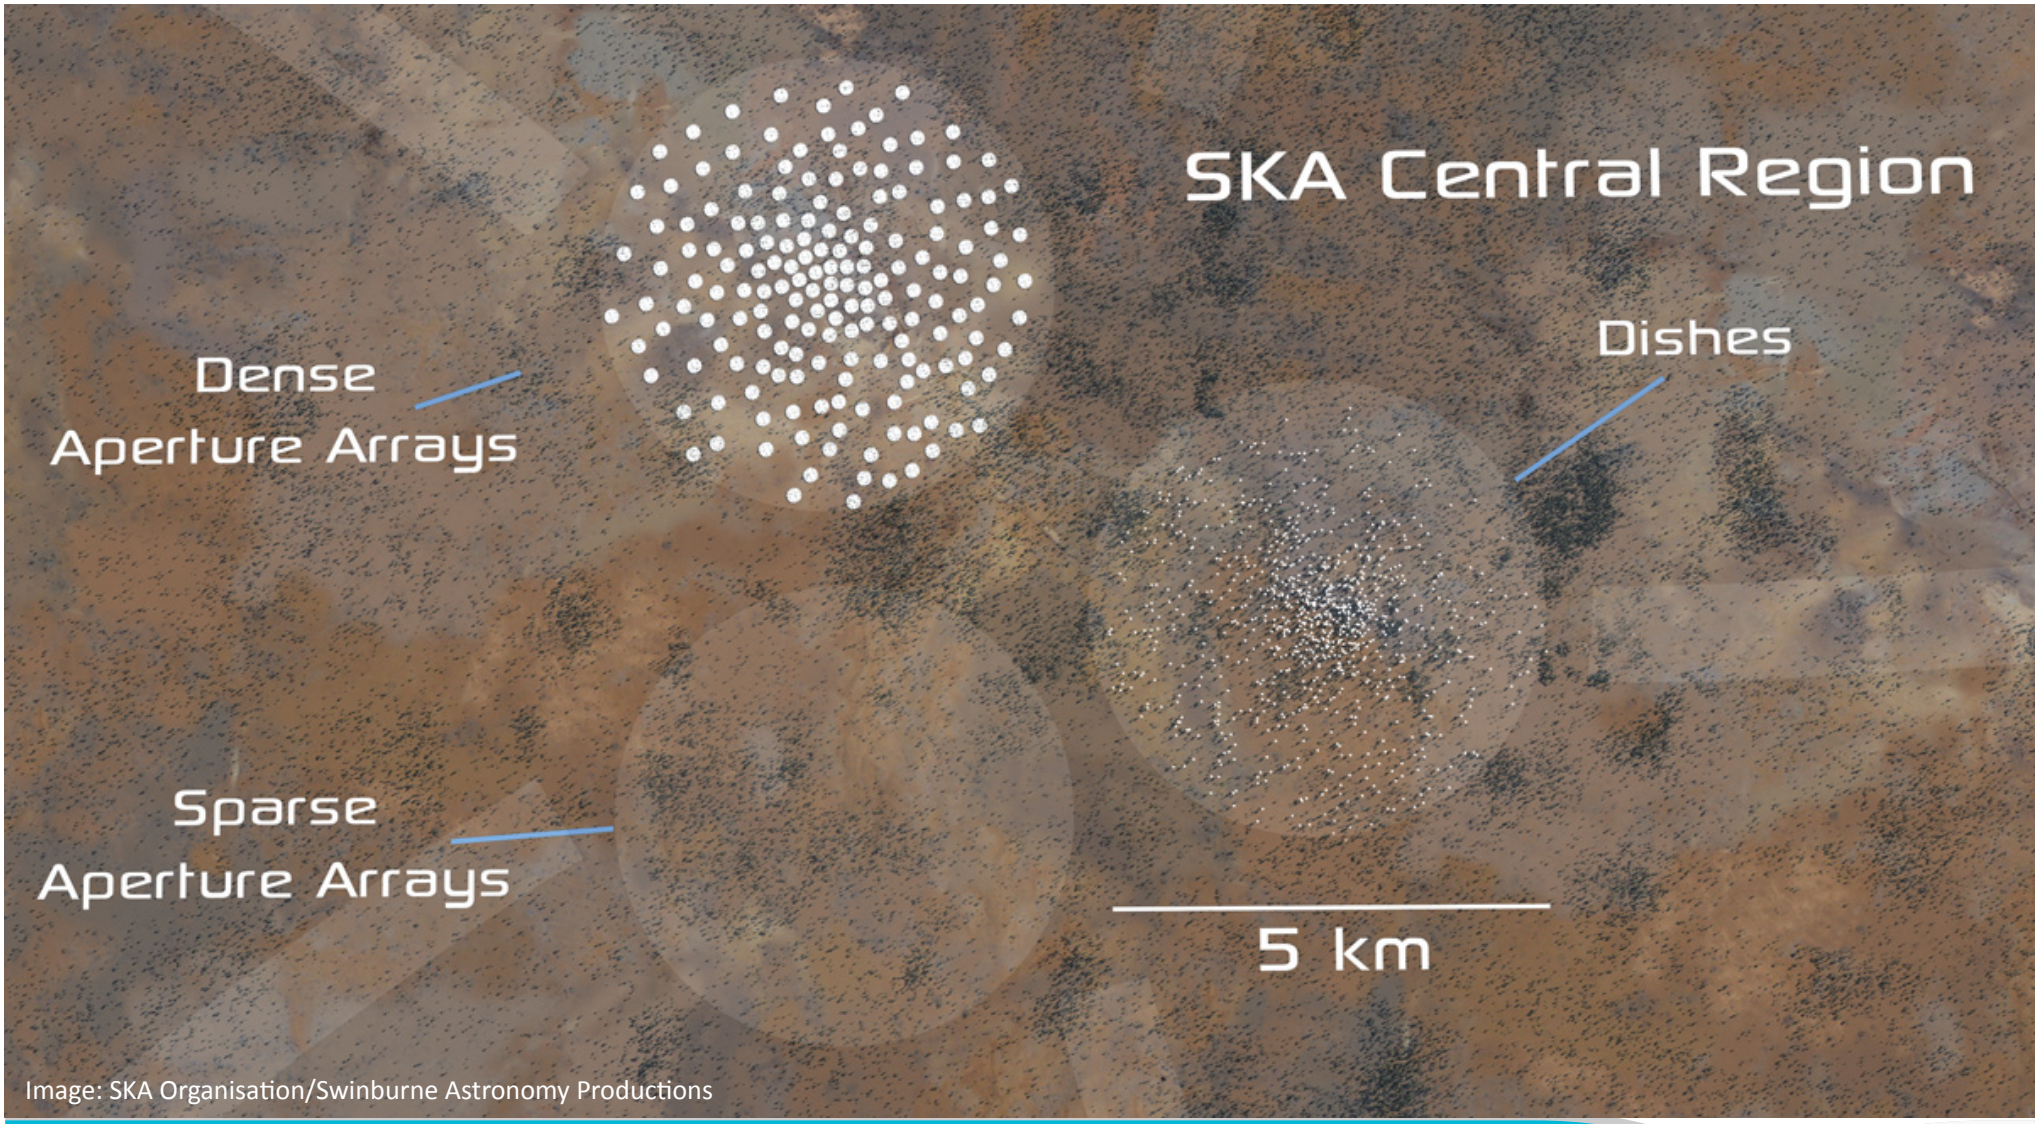

Image: SKA Organisation/Swinburne Astronomy Productions

SKA Dish Array

Image: SKA Organisation/Swinburne Astronomy Productions

SKA Sparse Aperture Array

Image: SKA Organisation/Swinburne Astronomy Productions

SKA Dense Aperture Array

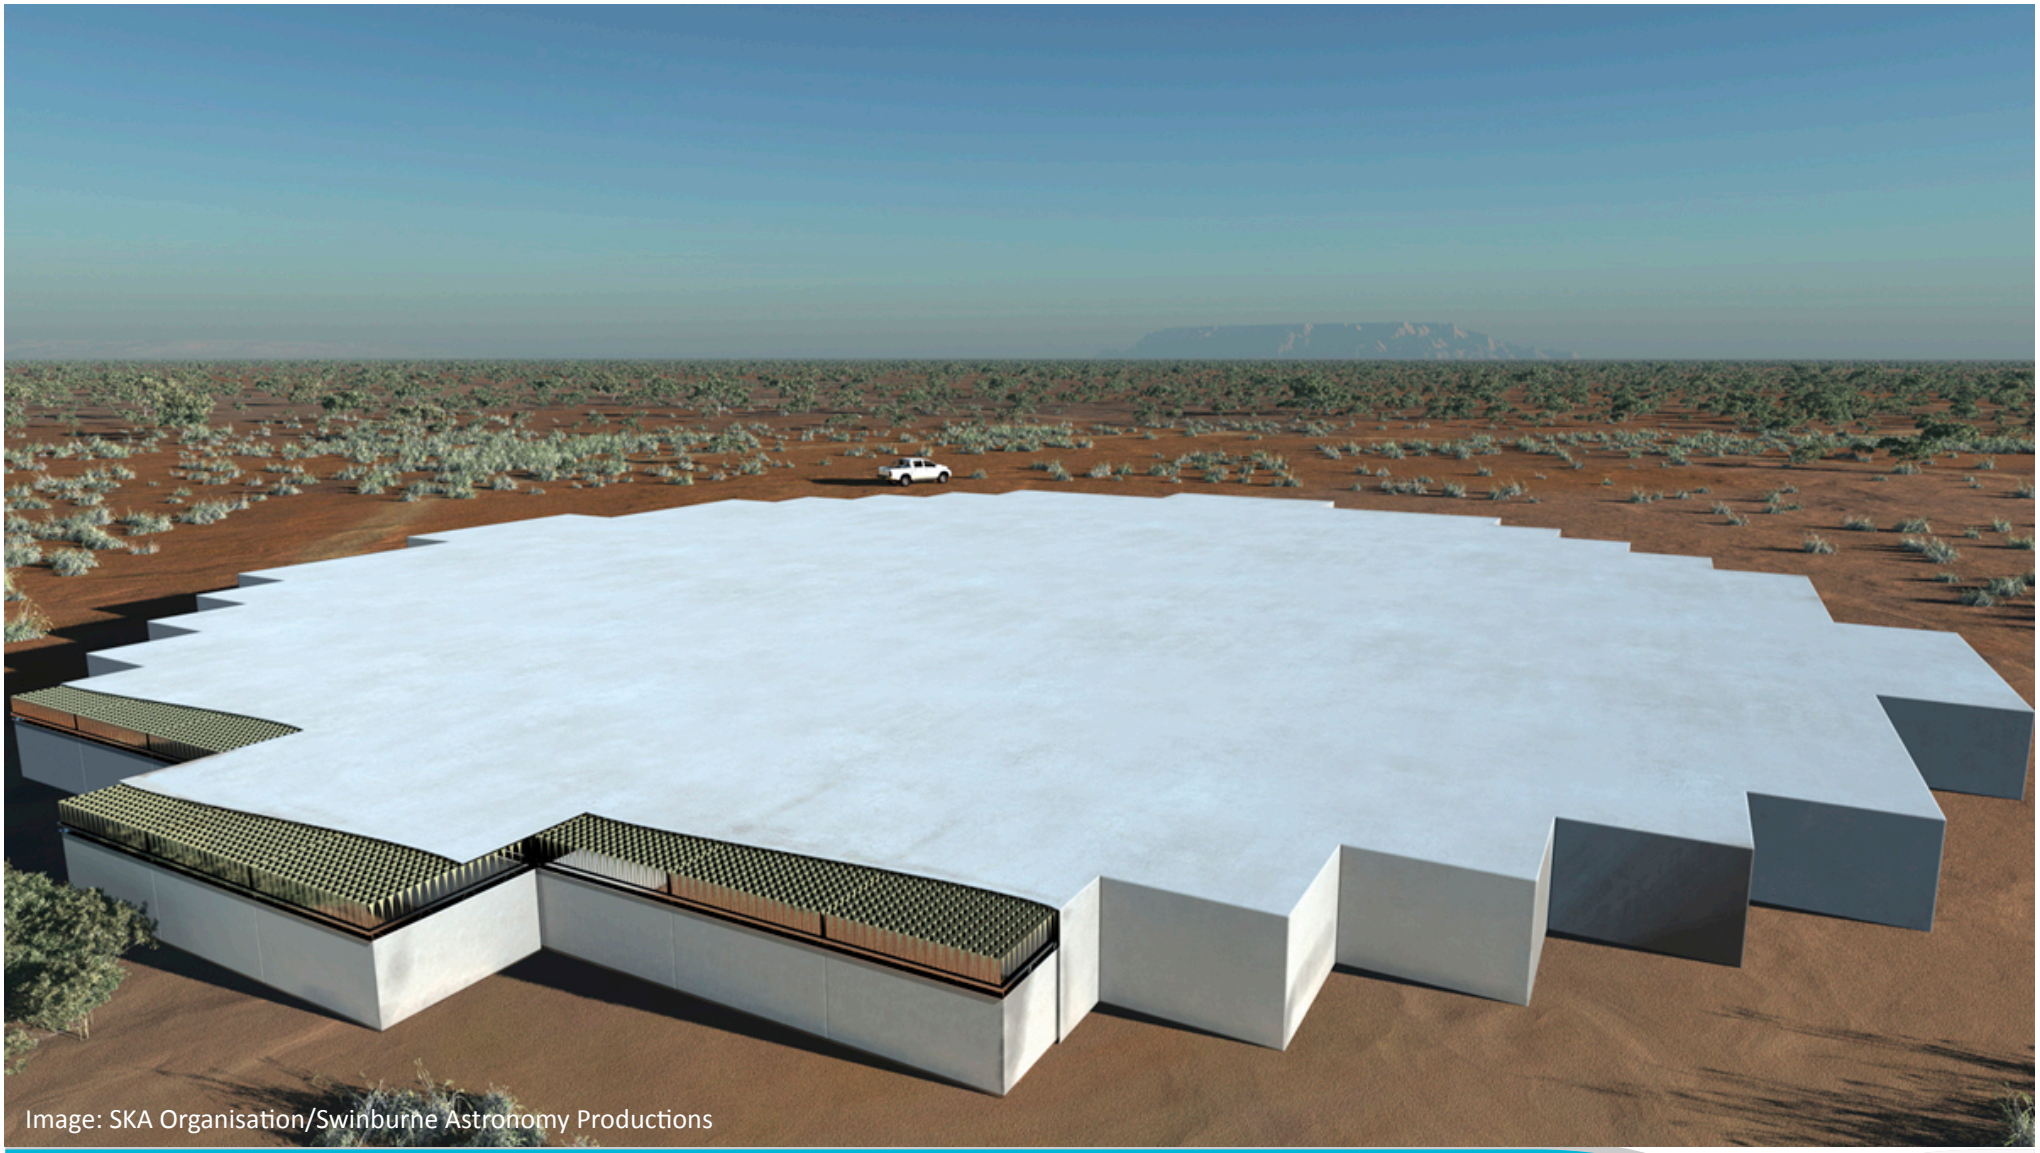

Image: SKA Organisation/Swinburne Astronomy Productions

Project Timeline

- **2006** Short listing of suitable sites
- **2008-12** Conceptual design
- **2012** Site selection
- **2013-15** Detailed design and pre-construction phase
- **2016-20** Phase one (SKA1) construction
- **2020-24** Phase two (SKA2) construction

Site Decision – SKA Phase 1 (SKA1)

- **SKA1_LOW (0.07 – 0.45GHz)**

  50 stations of low frequency aperture array antennas, with approximately 10,000 antennas per station, will be located in **Australia**

- **SKA1_MID (0.45 – 3GHz)**

  190 SKA dishes and 64 MeerKAT dishes equipped with single pixel feeds will be located in **South Africa**

- **SKA1_SURVEY (0.45 – 3GHz)**

  60 SKA dishes and 36 ASKAP dishes equipped with phased array feeds will be located in **Australia**

Site Decision – SKA Phase 2 (SKA2)

- **SKA2_LOW (0.07 – 0.45GHz)**

  The low frequency aperture arrays will be extended to 250 stations in **Australia**

- **SKA2_MID (0.45 – 10GHz)**

  The dish array will be extended to about 3,000 dishes with a maximum baseline of 3,000 km across **Southern Africa**

- **SKA2_AA – South Africa (0.4 – 1.4GHz)**

  A new component comprising 250 mid frequency aperture array stations will be located in **South Africa**

SKA2 Data Rates

SKA2_LOW – 250 aperture array stations (Aus)

250 stations × (380×10^6) samples per second × 2 (Nyquist) × 480 beams × 2 pols × 8-bits per sample

= **1.3 Petabits per second**

SKA2_MID – 3000 dishes with single pixel feeds (SA)

3000 dishes × 10^9 samples per second × 2 (Nyquist) × 2 pols × 8 bits per sample

= **87 Terabits per second**

Note: Figures are estimates only

Image Cubes

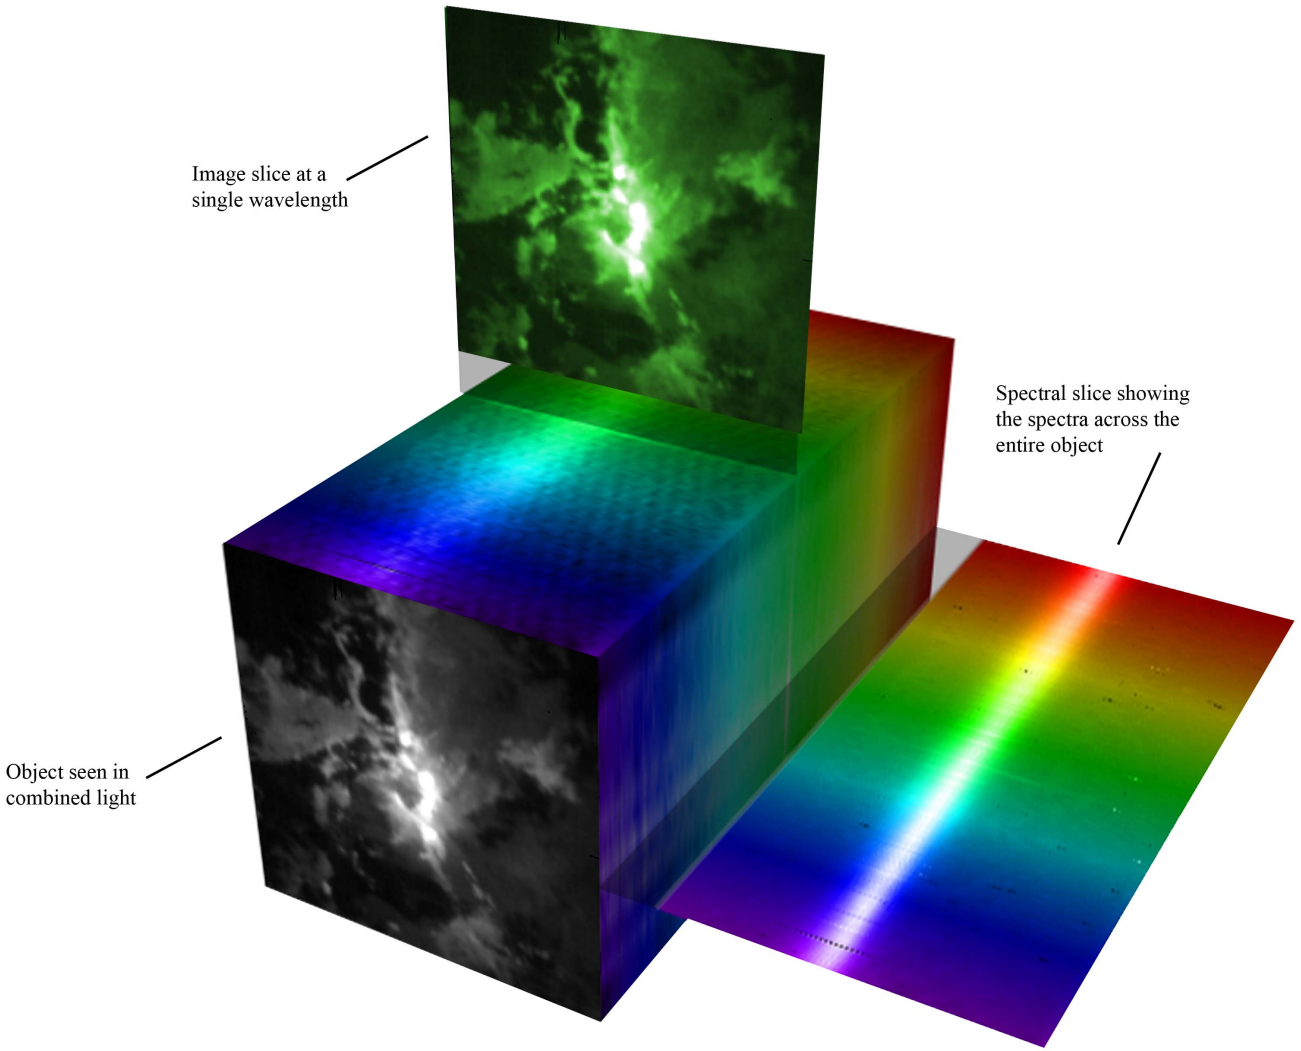

Image: Joint Astronomy Centre, The University of Hawaii

Image Cubes

SKA2_LOW – 250 aperture array stations (Aus)

120,000 (RA) × 120,000 (Dec) × 38,000 (Spectral Channel) × 32-bits per pixel

= **1.9 Petabytes**

SKA2_MID – 3000 dishes with single pixel feeds (SA)

80,000 (RA) × 80,000 (Dec) × 38,000 (Spectral Channel) × 4 polarisations × 32-bits per pixel

= **0.86 Petabytes**

Note: Figures are estimates only

Climbing Mount Exaflop

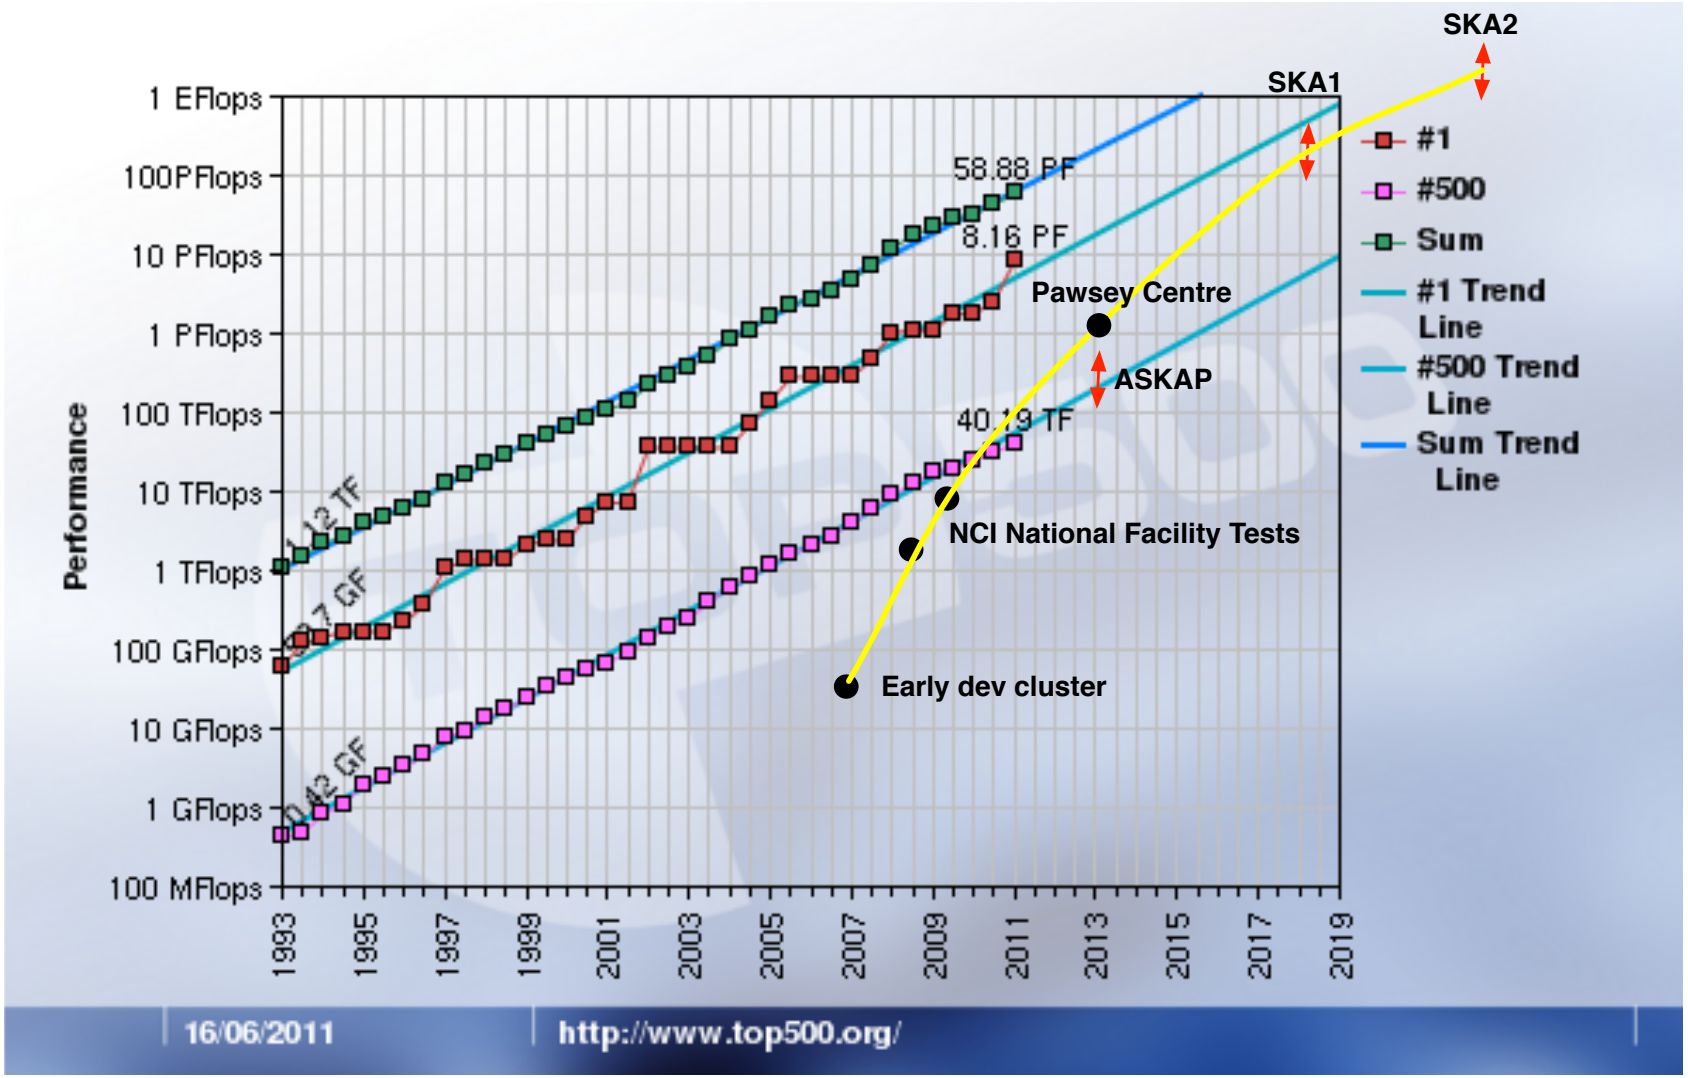

Lessons Learnt and Challenges Ahead

(Networking) Lessons Learnt - 1

- Transceiver and connector “hell”:
  - 1Gbit/s: GBIC, SFP: world has converged on SFP,
  - 10Gbit/s: XENPAK, XFP, X2, SFP+, …: not all media types available in all physical packages (e.g. LX4 and LRM),
  - ST, FC, SC, SCA, LC, LCA, E2000, SMA, … : minimise if possible.

Minimise transceiver types (depends on equipment choice) and connector type where possible

Where copper needs to be used, avoid CX4 connectors and use Cat6 or better cabling

(Networking) Lessons Learnt - 2

- Fibre types:
  - Multimode:
    - OM1 and OM3 common,
    - OM4 has now been standardised.
  - Singlemode:
    - patch leads and the like much easier,
    - active equipment more expensive (LASERs over VCSEL or LEDs).

Use single-mode only, even for short-haul data-centre links

(Networking) Lessons Learnt - 3

- In-equipment Optical Power Monitoring:
  - transmission equipment usually very good,
  - LAN equipment typically poor but improving (actually it is better supported in 10Gbit/s hardware than 1Gbit/s),
  - transceivers need to support this as well.

Buy transceivers that support "DOM" and before purchase check if active equipment supports this feature

(Networking) Lessons Learnt - 4

- **Take advantage of what the telecommunications industry have learnt over many decades:**
  - civil works,
  - CEVs,
  - good quality “as-built” documentation,
  - test results.
  - independent testing and verification is important.
- **Avoid media-converters!**
- **Understand the performance and limitations of active equipment prior to purchase (be annoying and pester the vendors until you get an answer you understand).**
- **Have a network architecture and standards document.**

SKA Networking Issues

- Topology, terminations, fibre plant management.
- Classes of network:
  - “production”,
  - monitor and control,
  - science data,
  - safety critical functions,
  - …
- Redundancy
- Protocols and interfaces
- Bandwidth:
  - Local,
  - National,
  - International.
- Cost

Remember…

The Network is the Telescope

The design cannot be "cast in stone" from Day 1. It is scientific research which often pushes hardware and software to their limits.

Thank you

CSIRO Astronomy and Space Science
Dr Shaun W Amy
Data Transmission Specialist/Computing Infrastructure Group Leader

t +61 2 9372 4452
e Shaun.Amy@csiro.au
w http://www.atnf.csiro.au