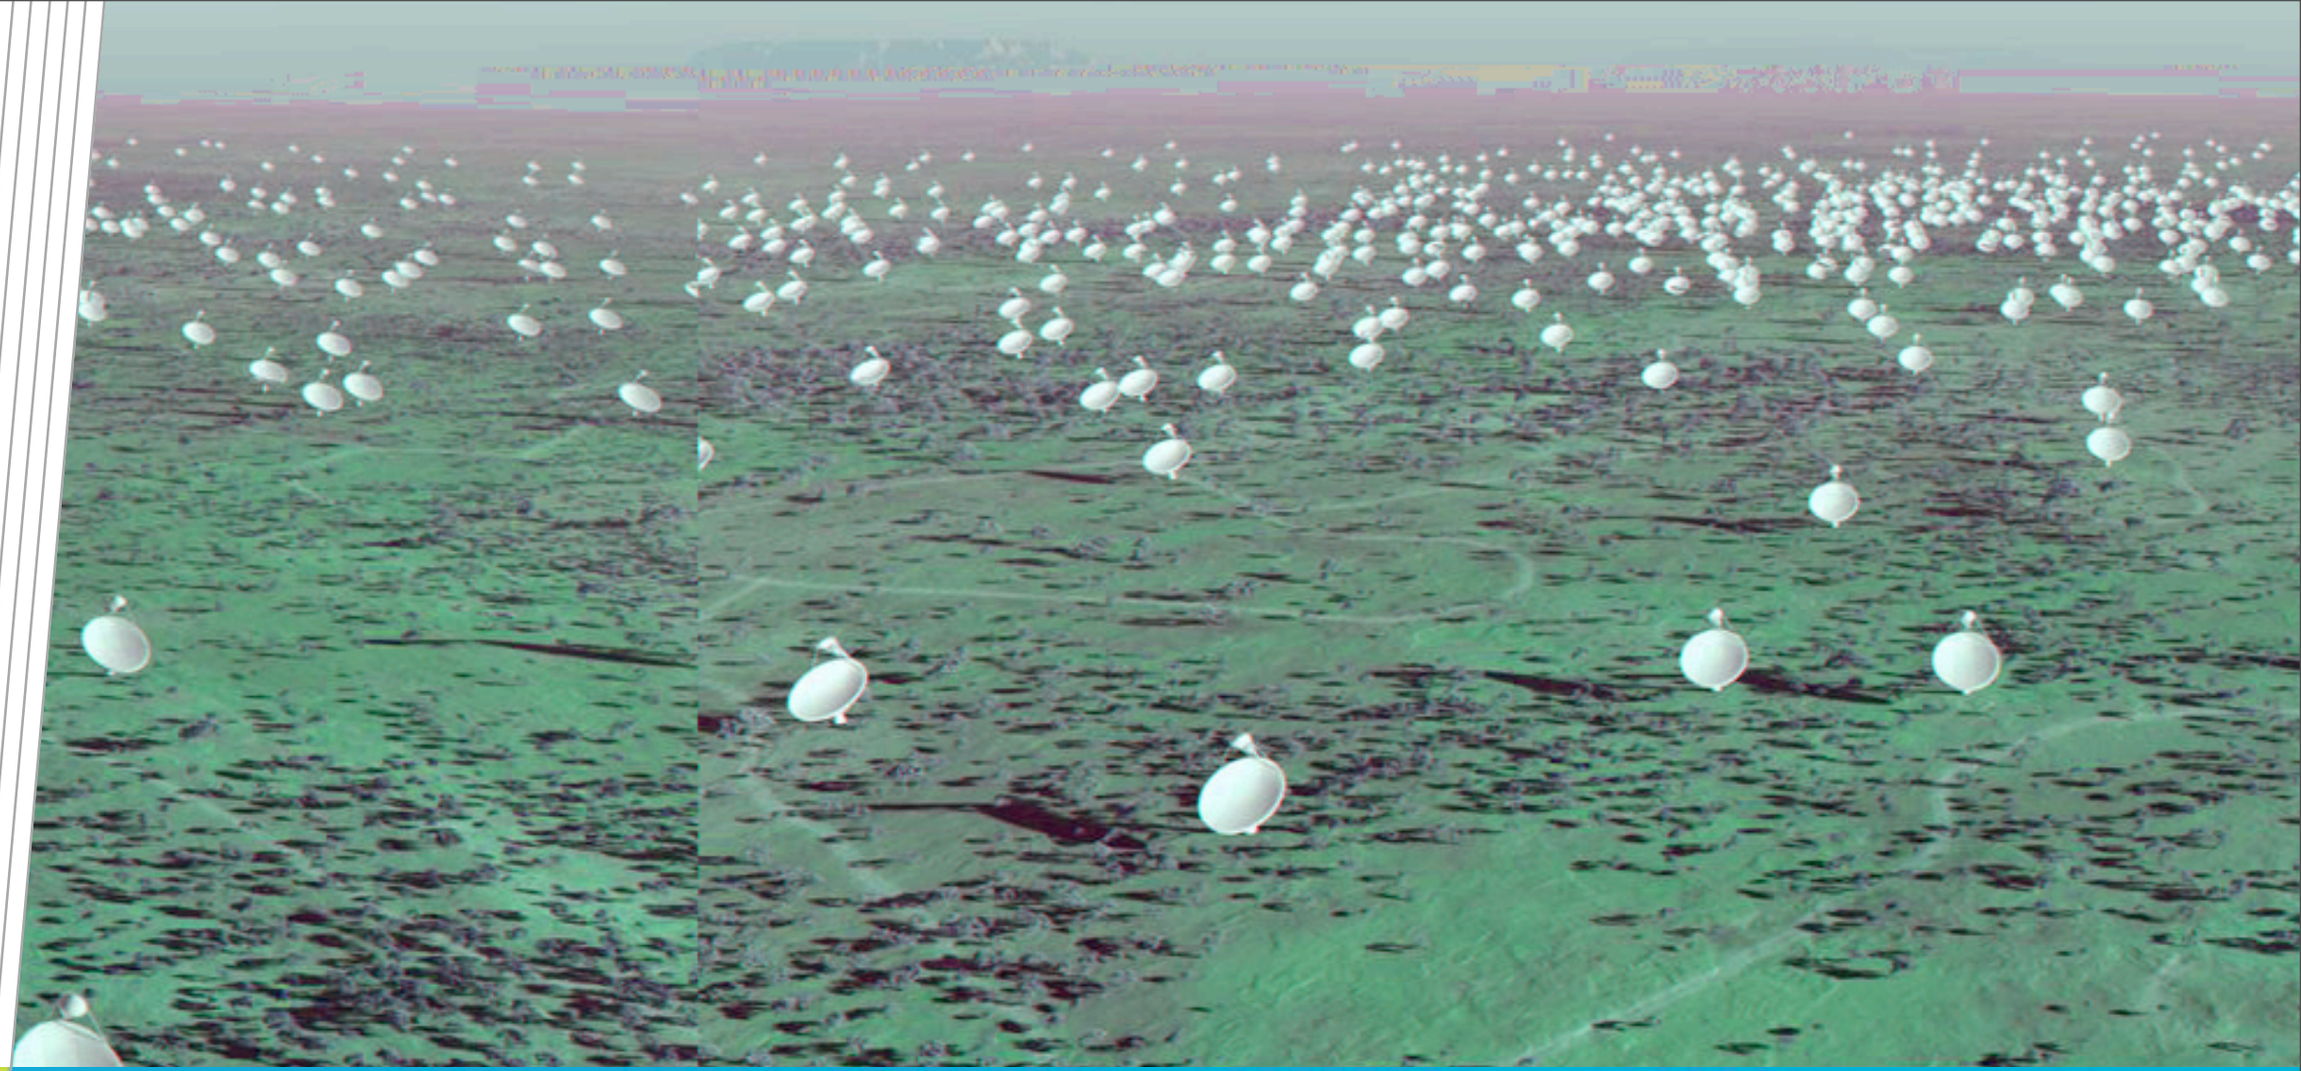

www.csiro.au

Square Kilometre Array as a Co-Design Vehicle

Tim Cornwell, ASKAP Computing Project Lead
Ben Humphreys, ASKAP Computing Project Engineer
Australian Square Kilometre Array Pathfinder

CSIRO

SKA Science

- Key Science Projects for the SKA
  - Hydrogen survey – dark energy
  - Pulsar survey – strong field tests of gravity
  - Cosmic magnetism – origin of B fields
  - Cradle of Life
  - Epoch of Re-ionization
  - Exploration of the Unknown

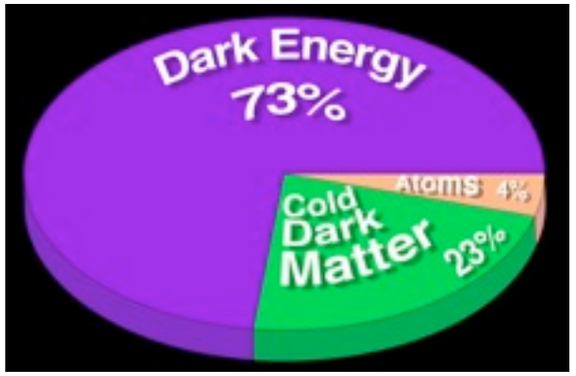

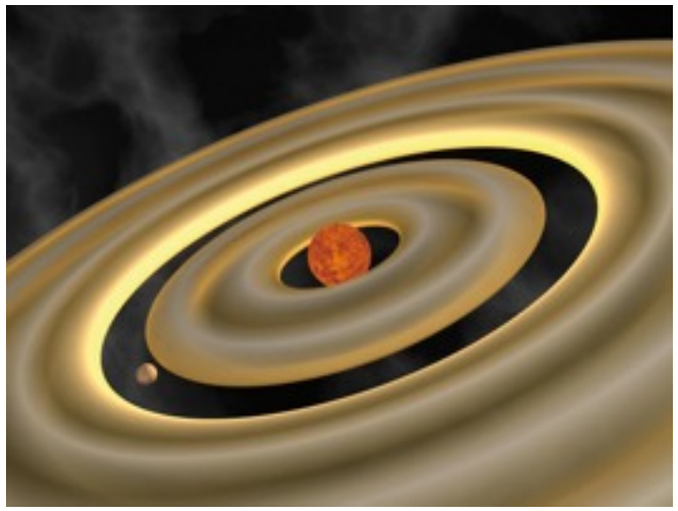

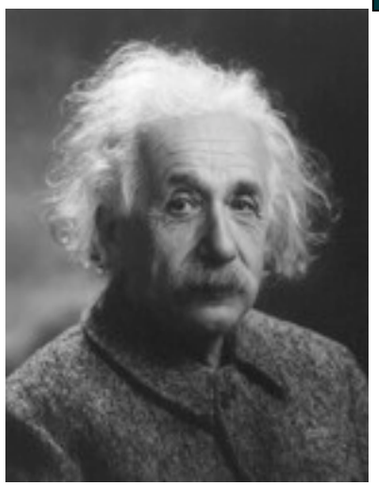

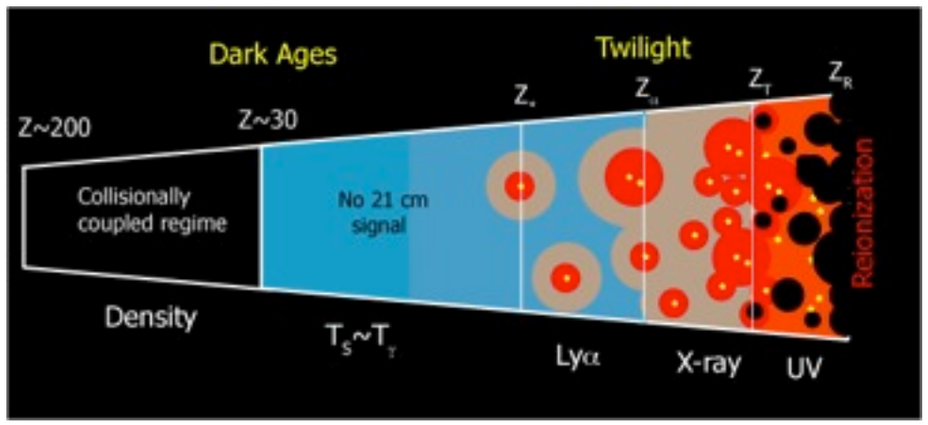

The Square Kilometre Array

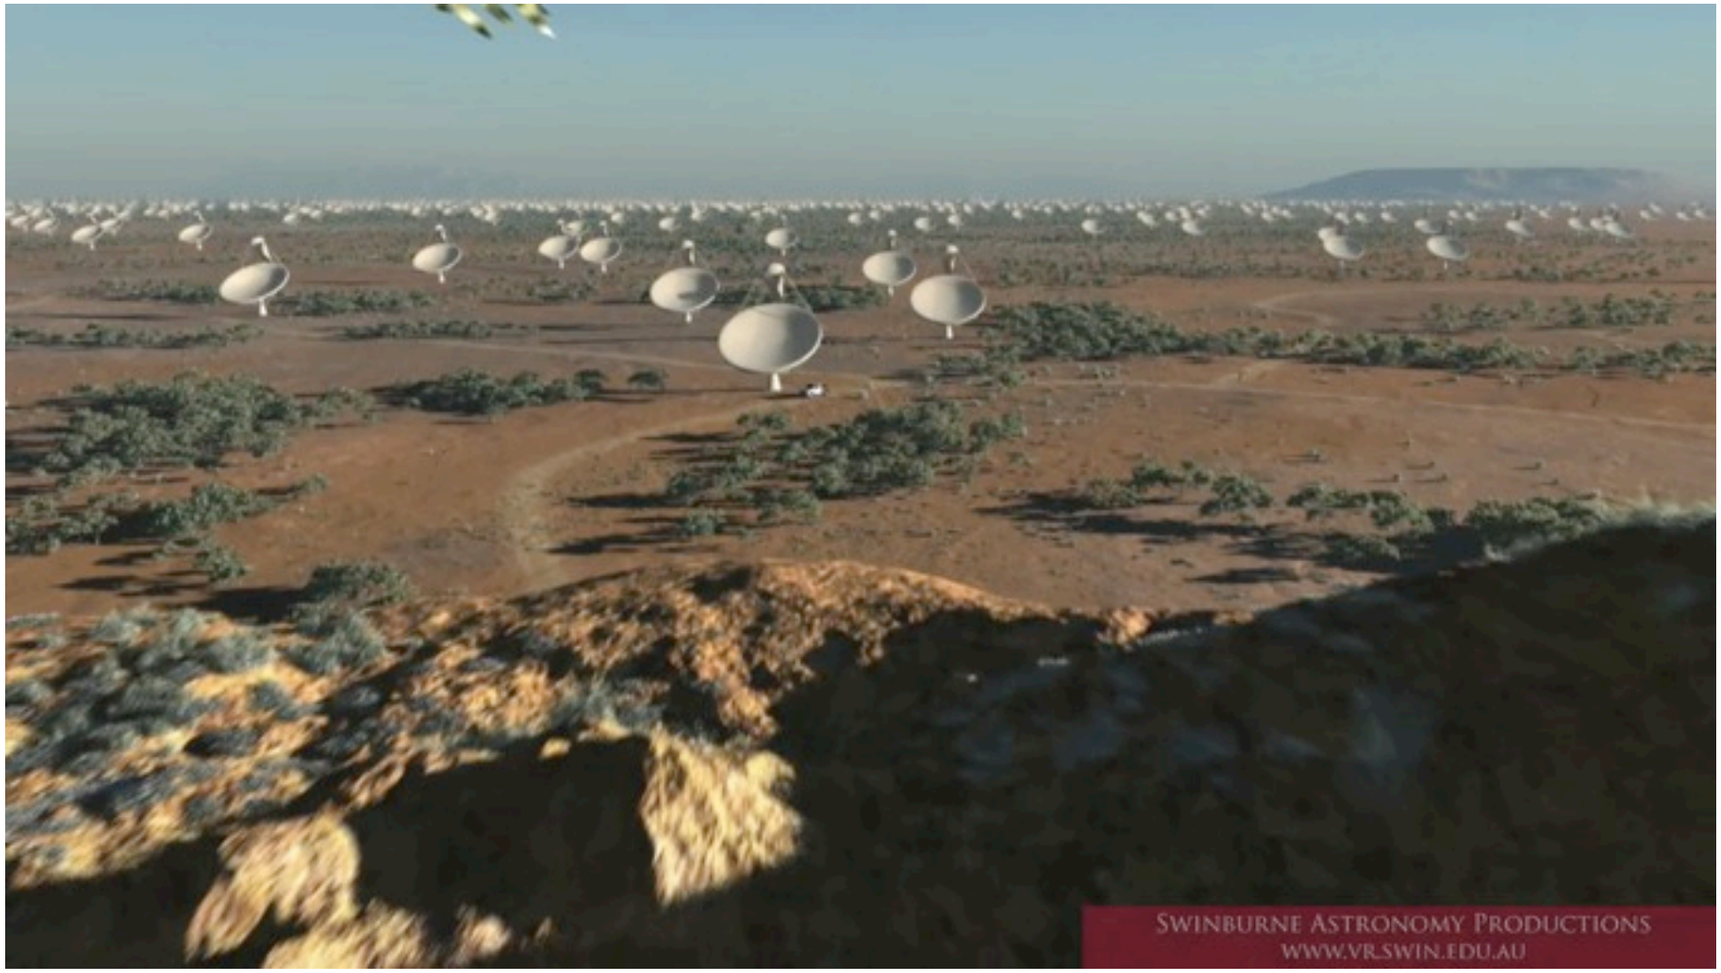

- 2020 era radio telescope
- Very large collecting area (km^2)
- Very large field of view
- Wide frequency range (70MHz - 25 GHz)
- Large physical extent (3000+ km)

- International project
- Telescope sited in Australia or South Africa
- Headquarters in Europe or US
- Multiple pathfinders and precursors now being built around the world

LOFAR

- Constructed by ASTRON in the Netherlands
- **L**ow **F**requency **AR**ray
- Sparse Aperture Array
- 10 - 80 MHz, 120 - 240 MHz
- Correlator implemented in Blue Gene/L/P

- Opened June 12 2010

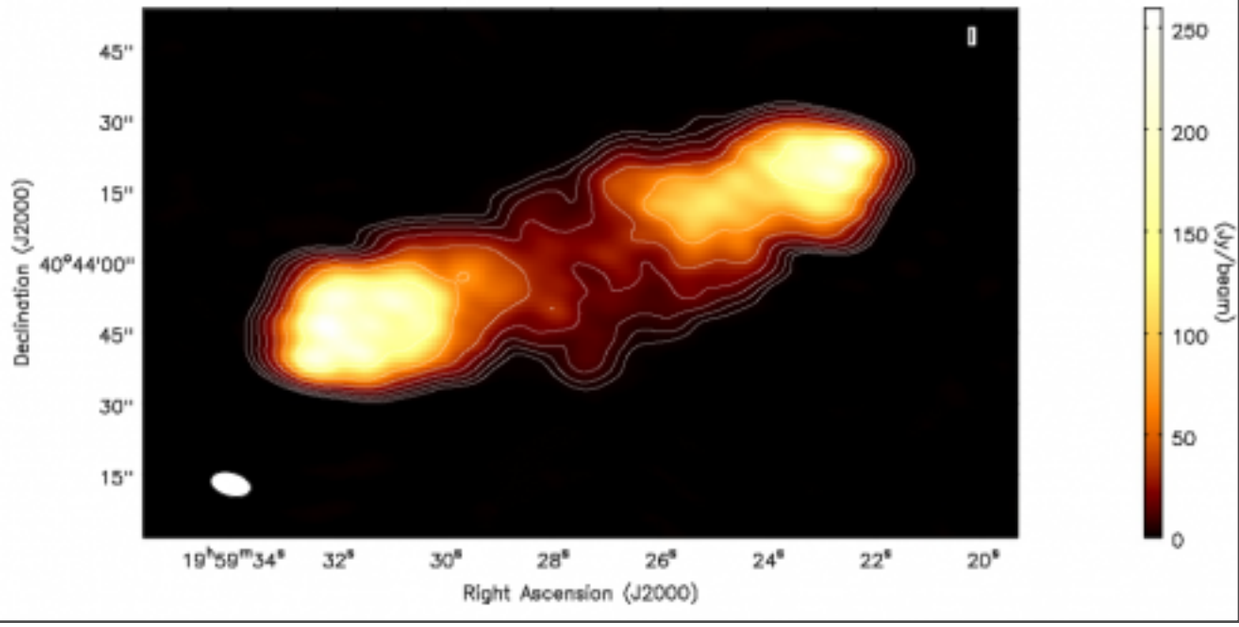

USA SKA Technology Development Program

- http://skatdp.astro.cornell.edu/
- Computer Processing Group
- Concentration on calibration and imaging data processing
- Mathematical analysis of processing methods

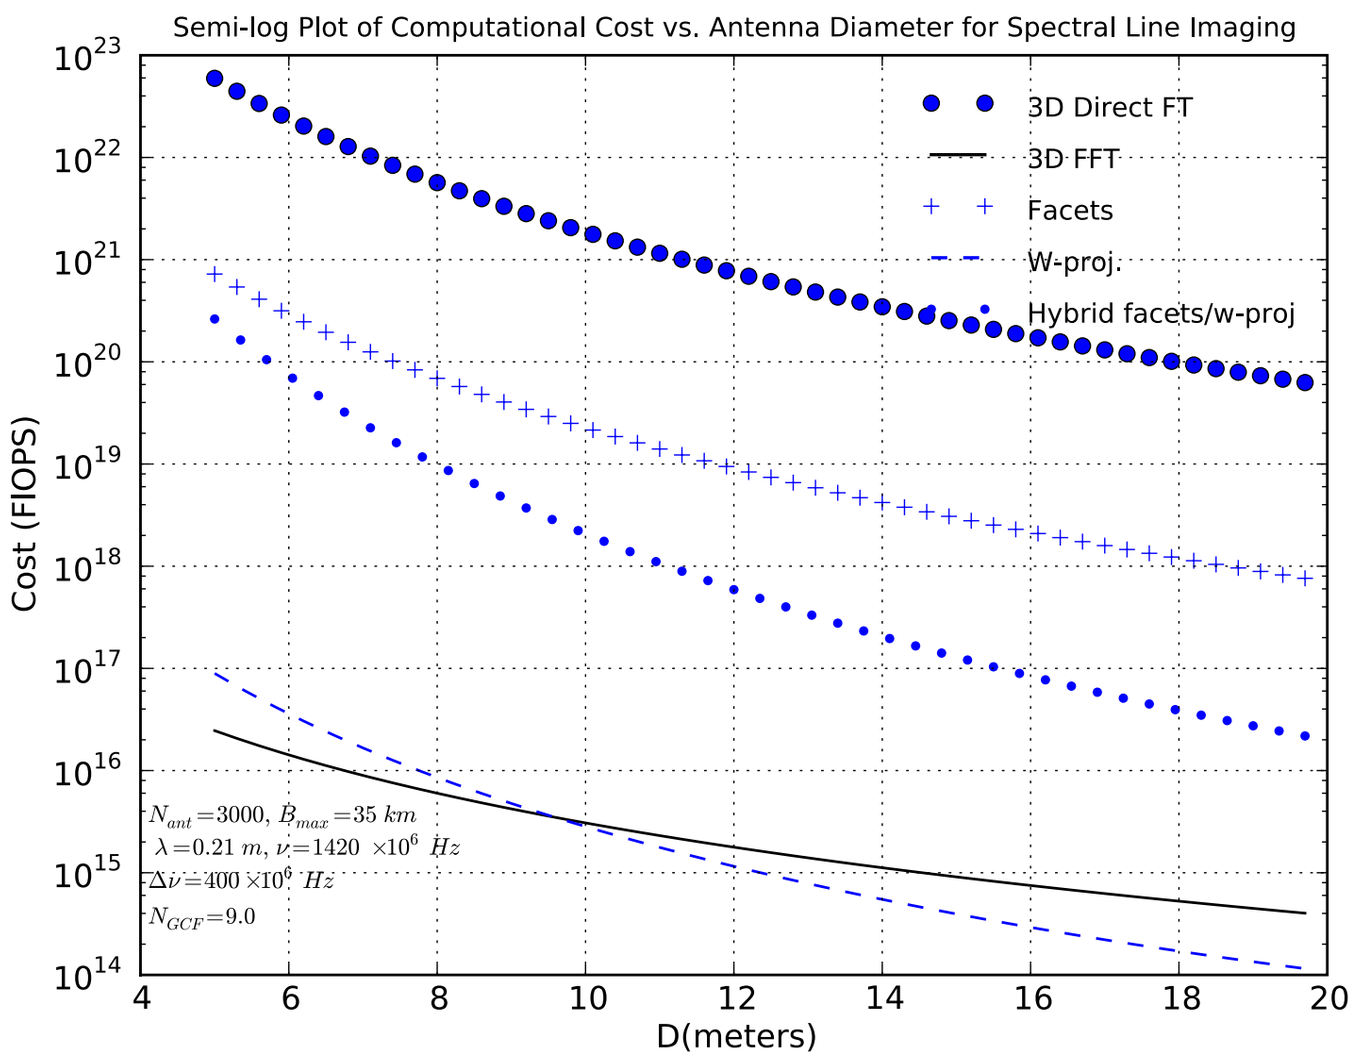

Australian SKA Pathfinder = 1% SKA

- Fastest survey radio telescope in the world
- Sited at Boolardy, Western Australia
- 36 antennas compared to ~ 3600 for SKA
- First 6 antennas installed

- 150MA$
- Early test observations 2011
- Full observing 2013
- Demonstrates wide field of view technology for SKA

CSIRO

Australian SKA Pathfinder = 1% SKA

- Fastest survey radio telescope in the world
- Sited at Boolardy, Western Australia
- 36 antennas compared to ~ 3600 for SKA
- First 6 antennas installed

- 150MA$
- Early test observations 2011
- Full observing 2013
- Demonstrates wide field of view technology for SKA

ASKAP data flow

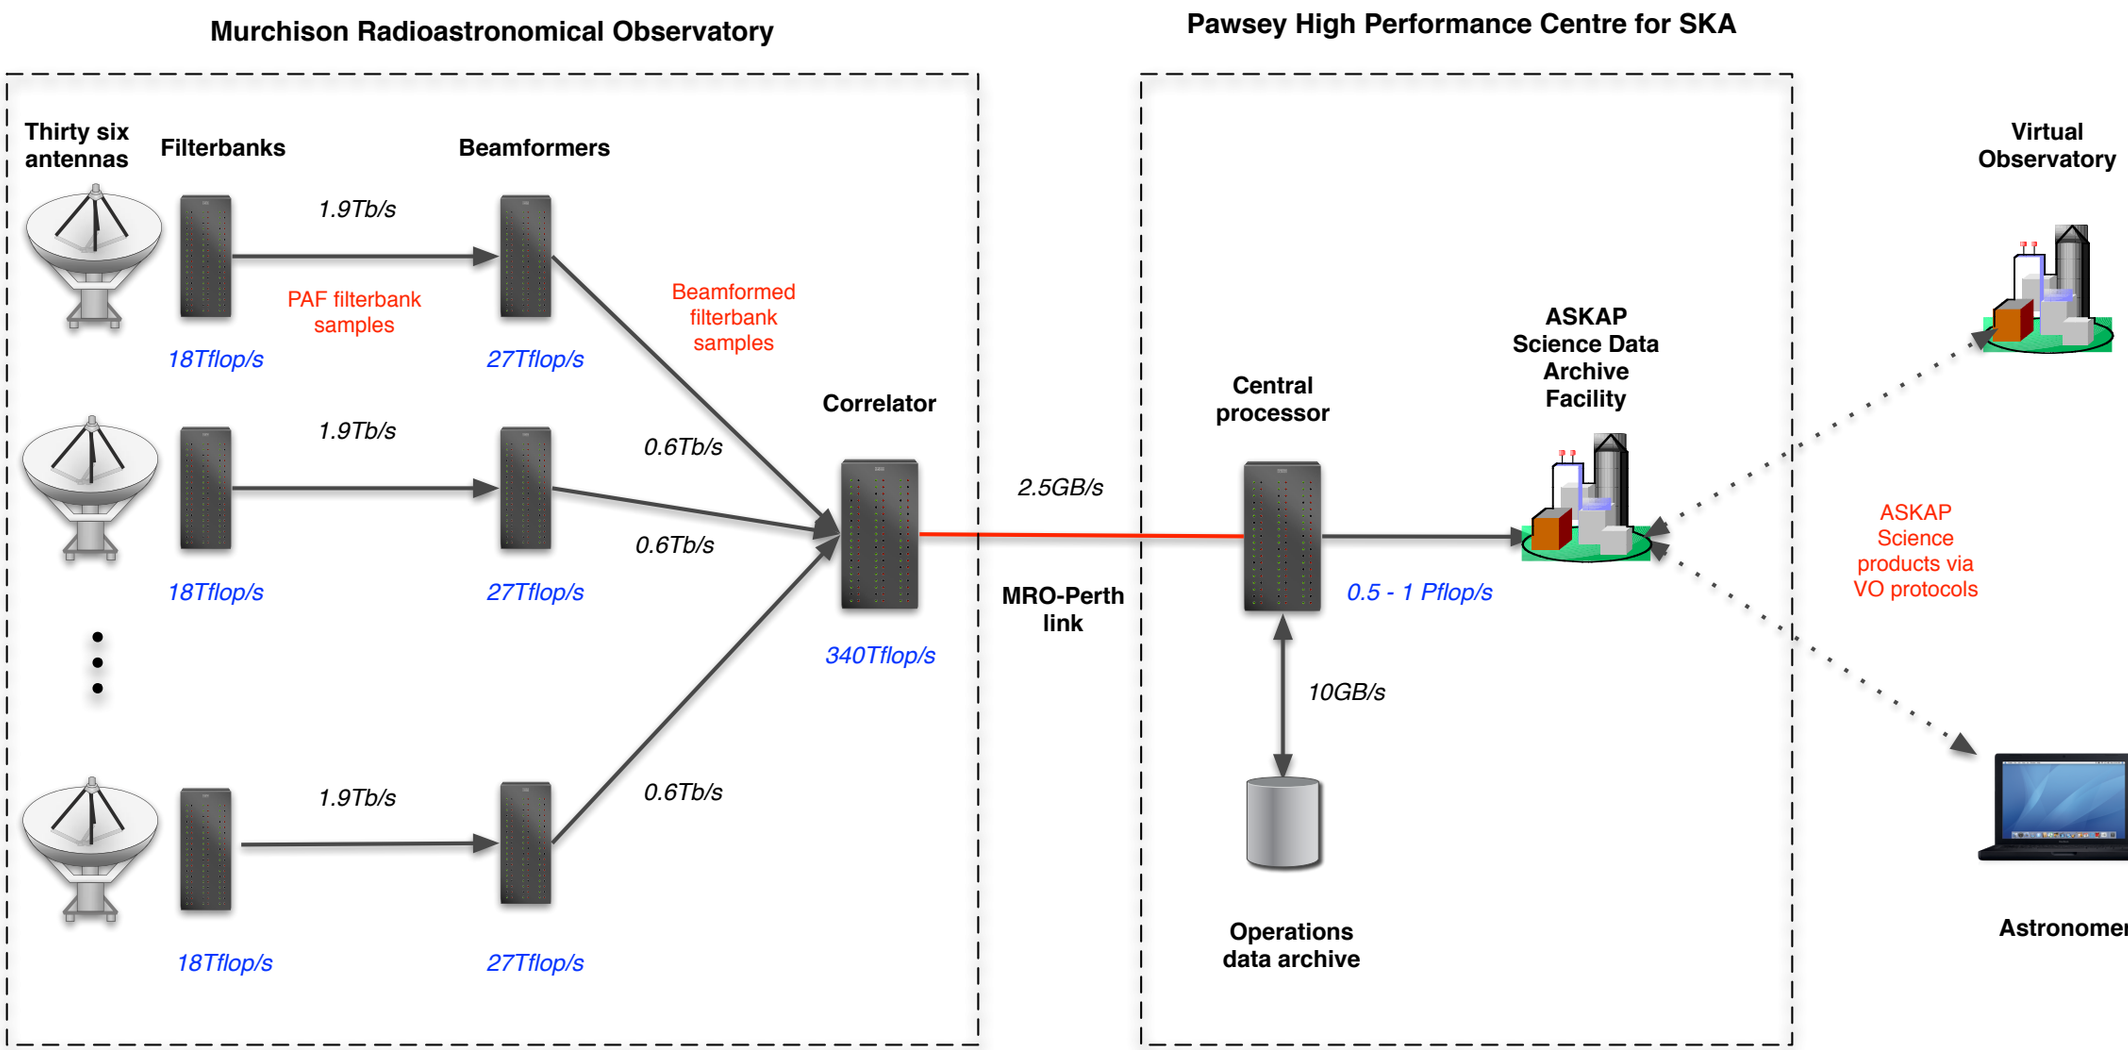

- From observing to archive with no human decision making
  - Calibrate automatically
  - Image automatically
  - Form science oriented catalogues automatically

T. Cornwell, July 9 2010

Pawsey High Performance Computing Centre for SKA Science, Perth, Western Australia

- A$80M, funded by Australian Federal government
- 100 TF machine being commissioned
  - HP cluster in a box
  - Full scale ASKAP test during commissioning
  - Used in 2011 for early telescope testing
- Petascale system by 2013
  - 25% for radio astronomy
- Would like to access to 1PF before that for scaling!

CSIRO

iVEC THE PAWSEY CENTRE

Call for expressions of interest for early adopter use of the Pawsey Stage 1A system

iVEC is currently in the process of installing Pawsey Stage 1A, a Hewlett-Packard POD based system that will provide access to 9,600 cores based at iVEC@Murdoch.

This system will be commissioned in the next few months and iVEC is seeking expressions of interest from the iVEC user community for early adopter use during the 4 week commissioning phase.
Access during this period is aimed at providing time on the system so users can experiment and aid iVEC in configuring and validating the system.

To be eligible as early adopters users must be able to nominate a project capable of:

- generating a high impact research outcome within the commissioning phase by taking advantage of the Pawsey Stage 1A compute capacity ; or
- wanting to investigate performance of an existing or new application as preparation for future work.

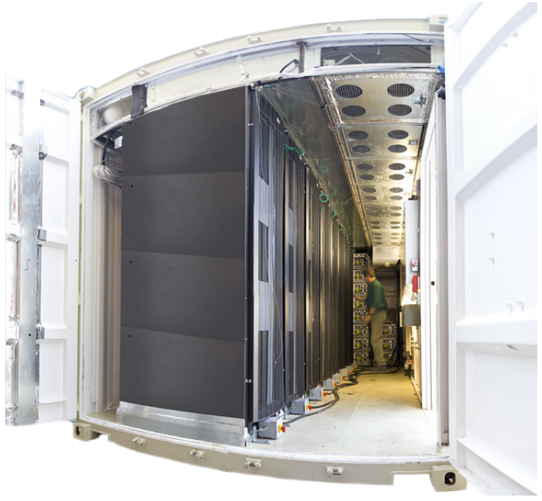

Preference will be given to projects that utilise thousands of cores .

Technical assistance from the iVEC staff will be available during the commissioning phase.

To register your interest please send a short description (less than one page) of your project and its requirements to **earlyadopters@ivec.org** , questions can also be sent to this address. **The deadline for expressions of interest is 22nd October, 2010** and we expect to announce the successful projects by the 5th of November.

After the commissioning phase, the Pawsey Stage 1A system will be available for more general access through the Pawsey Allocation Committee.

CSIRO

Radio interferometry

- Measured data = **Fourier/Fresnel transform** of sky

$$V_{ij,p} = \int \frac{I(l,m)}{\sqrt{1-l^2-m^2}} e^{j2\pi(ul+vm+w\sqrt{1-l^2-m^2})} \, dl \, dm$$

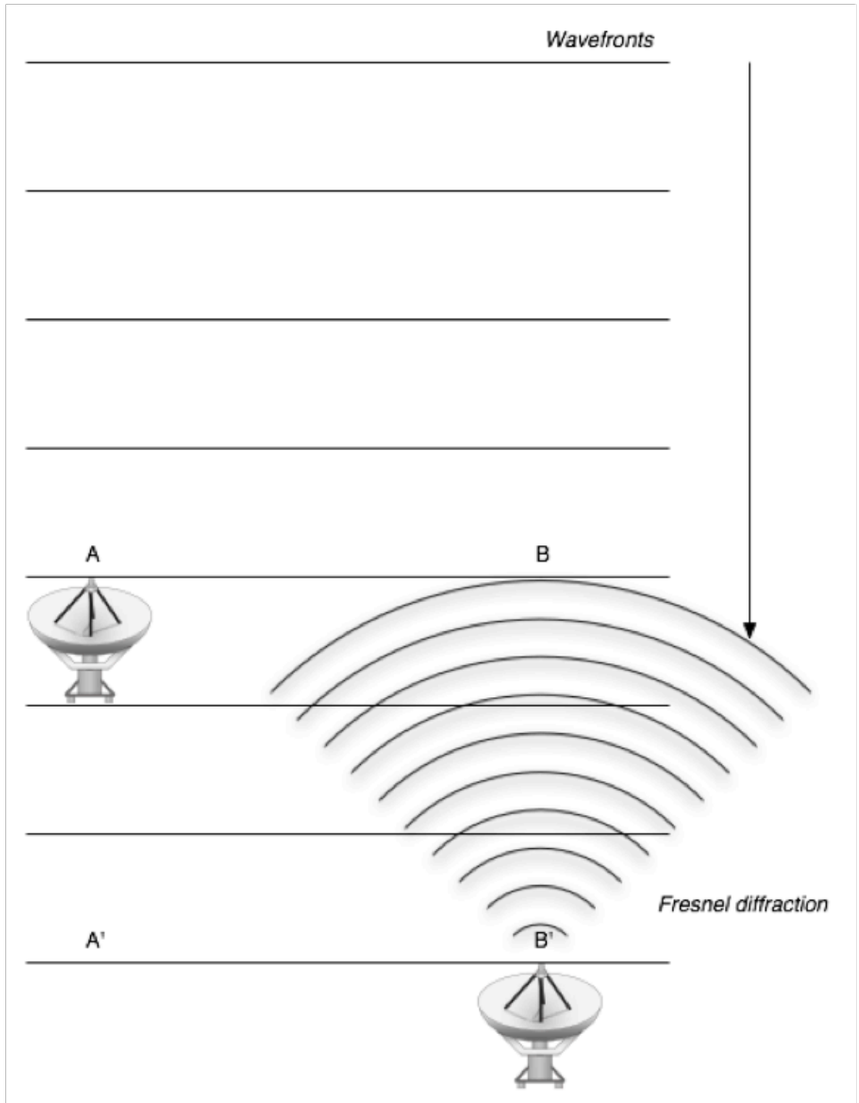

- Also must include directivity of receptors
- Summary paper:
  - Rau, U., Bhatnagar, S., Voronkov, M.A., Cornwell, T.J. (2009). *Advances in Calibration and Imaging Techniques in Radio Interferometry*. Proceedings of the IEEE, Vol 97 (8), 1472–1481.
- Intrinsically all-to-all or many-to-many
  - Depending on receptor type
- No closed form solution
- Iterative solutions exist
- Low arithmetic intensity in core algorithm
- SKA may allow single pass solution

(Some) Approaches to wide field imaging

Approach	How it works	Memory	CPU
Faceting	Divide sky into small patches for which the 2D transform is valid	Low	High
W stacking	Apply w-dependent phase screen by multiplication in image space	High	Low
W projection	Apply w-dependent phase screen by convolution in Fourier space	Low	High
Snapshot imaging	Partition by hour angle, zenith angle, FFT, interpolate to same image plane	Low	High
A projection	Apply antenna beams by convolution in Fourier space	Low	Low
AW projection	Antenna beam and phase screen by convolution in Fourier space	Low	High
A projection/W stack	Phase screen in image/antenna beam in Fourier space	High	Low

Co-design in SKA

- Kemball, Cornwell, and Yashar TDP CPG memo # 4

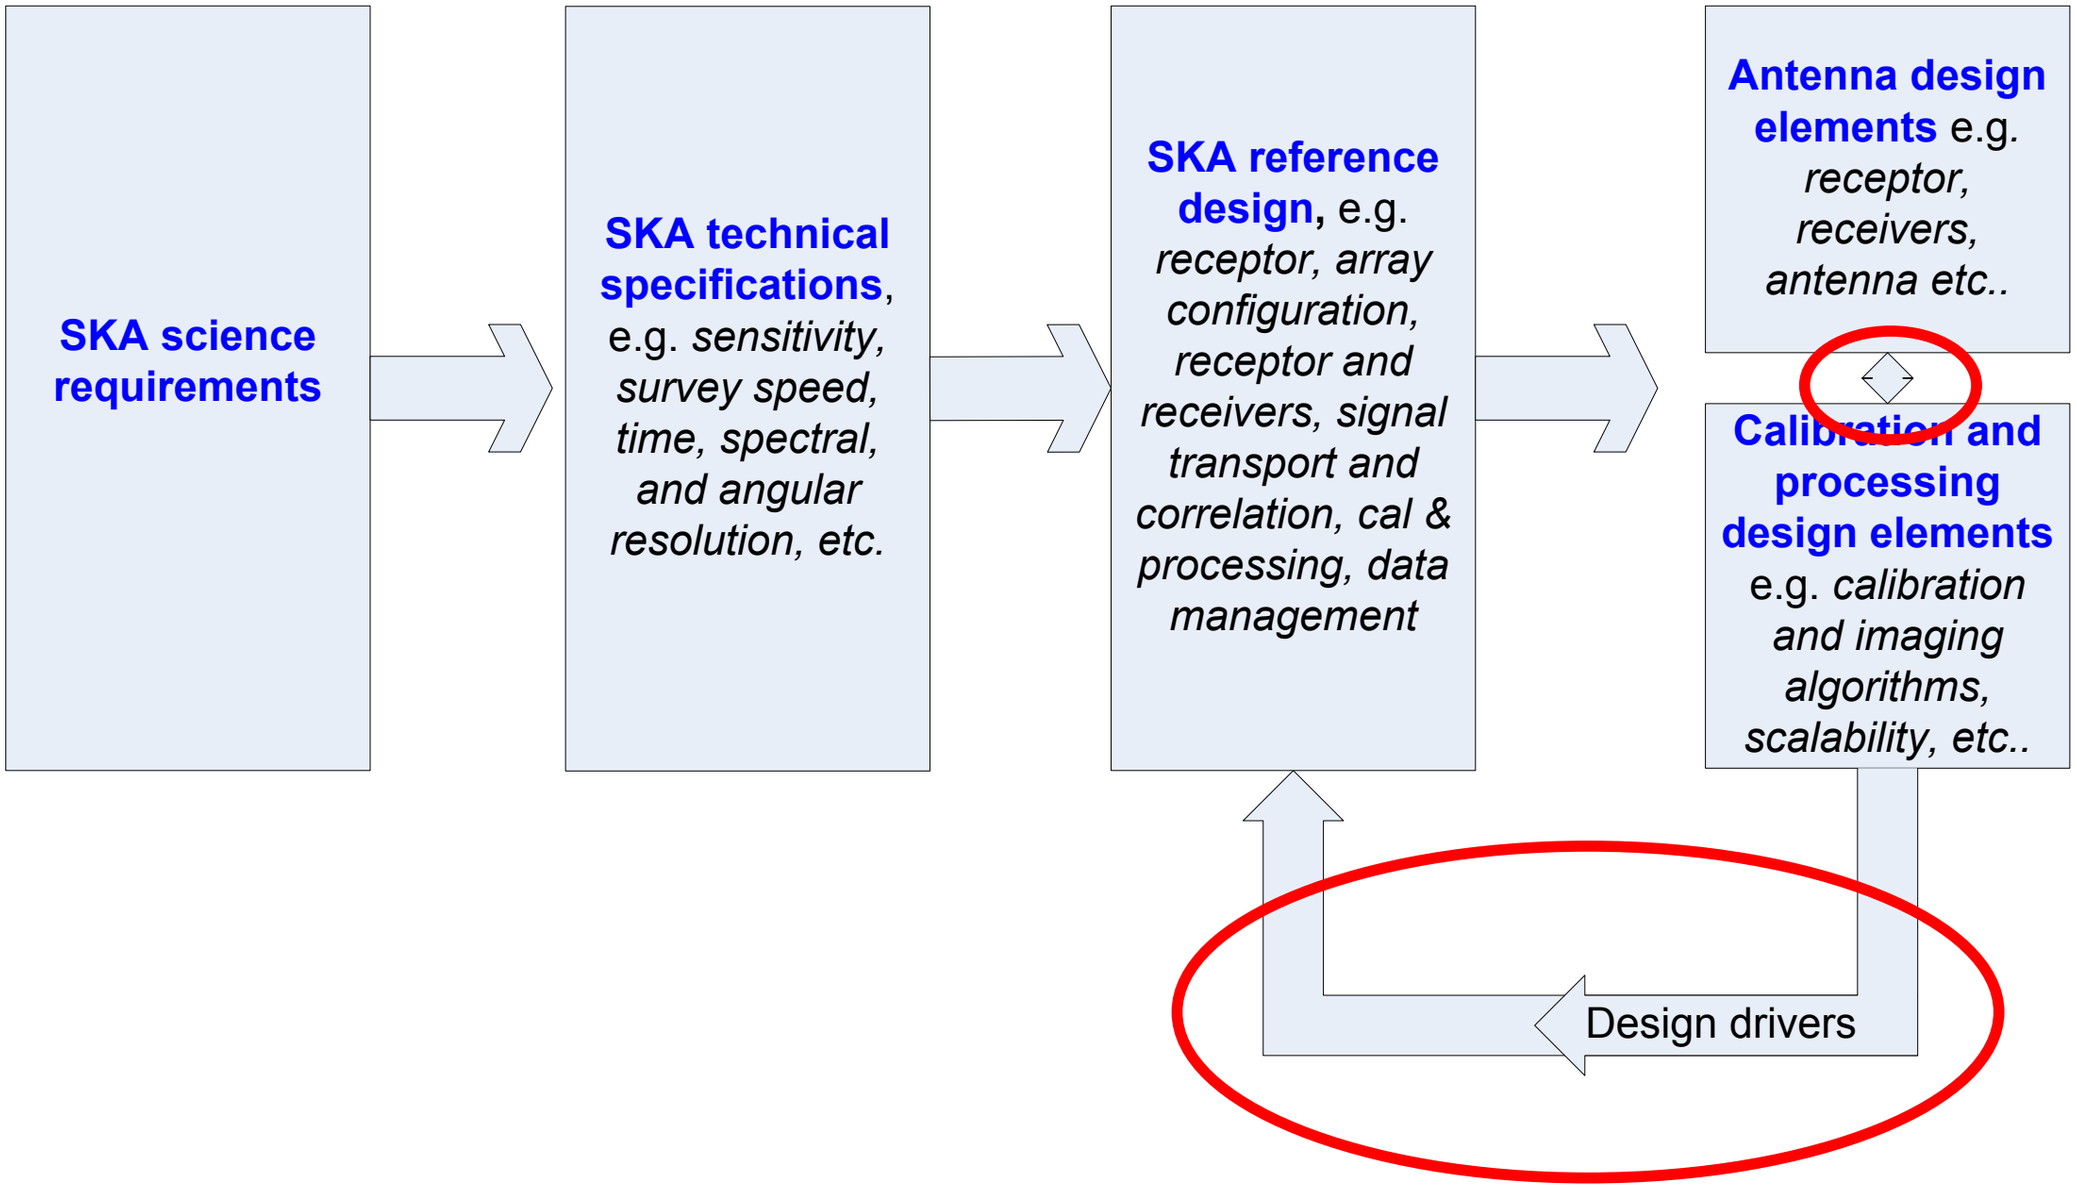

Telescope/software/algorithm co-design in ASKAP

- ASKAP has taken multiple steps to reduce computing costs and risk:

- ASKAP controls entire end-to-end software stack
- Focus on survey observations reduces complexity in software
- The phased array feed decouples radiation from different parts of sky - changes all-to-all to many-to-many
- The phased array feeds are kept fixed on the sky by a novel three-axis design for the antenna, thus obviating the need for software sky de-rotation and correction of scattering from the antenna focus supports.
- The phased array feed elements are calibrated at the front-end using radiators on the antenna surface.
- Astronomical calibration information is fed back to the telescope in real time, reducing some first order effects to second order.

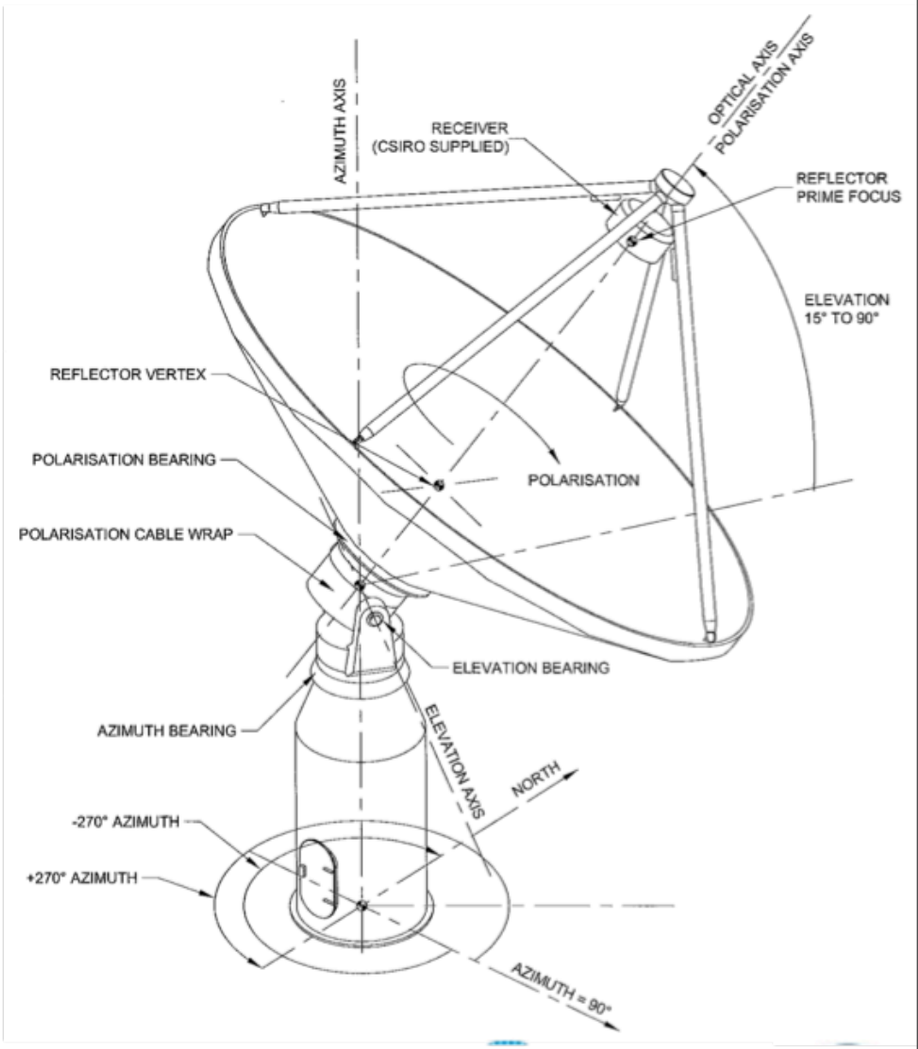

CSIRO

HPC challenges and lessons learned from ASKAP

- Streaming vs Batch processing
  - Telescope produces streams of data
  - No suitable battle-hardened, mature framework for handling high volume streaming data
  - Adopted traditional batch system controlled by DRMAA
  - Batch (or any buffered processing) definitely will not work for SKA
- Processor-memory gap
  - Older processors limited by memory latency
  - Newer processors can saturate memory bus
- Forced to shift from ~ 4GB/core to 1GB/core
  - ~ 2012 timescale
- Also important to investigate low-power systems GPU, BG/?
  - But (rate of gridding)/Watts ~ constant
- Vital to have multiple algorithms for same purpose

Comparison of ASKAP and SKA

	ASKAP	SKA
Observing model	Principally surveys	Surveys and targeted observations
Data per hour	10 TB	10 - 100 PB
Scale of processing	100TF – 1PF	10PF – 1EF
Type of parallelism	Iterative, almost Embarrassingly Parallel (EP) with occasional reconciliation	Single pass, almost EP with single reconciliation
Major costs	Convolutional resampling, iteration	Convolutional resampling (or better)
Arithmetic Intensity	Low	Low
Data flow nature and strategy	Intrinsically streamed, buffering to fast file system, barriers at reconcilation points	Intrinsically streamed, no buffering, no barriers, best try
Development and deployment environment	MPI/C++	????
Calibration	Multiple stages, iterative – requires data/image transforms	Multiple stages, iterative in image space only

Climbing Mount Exaflop

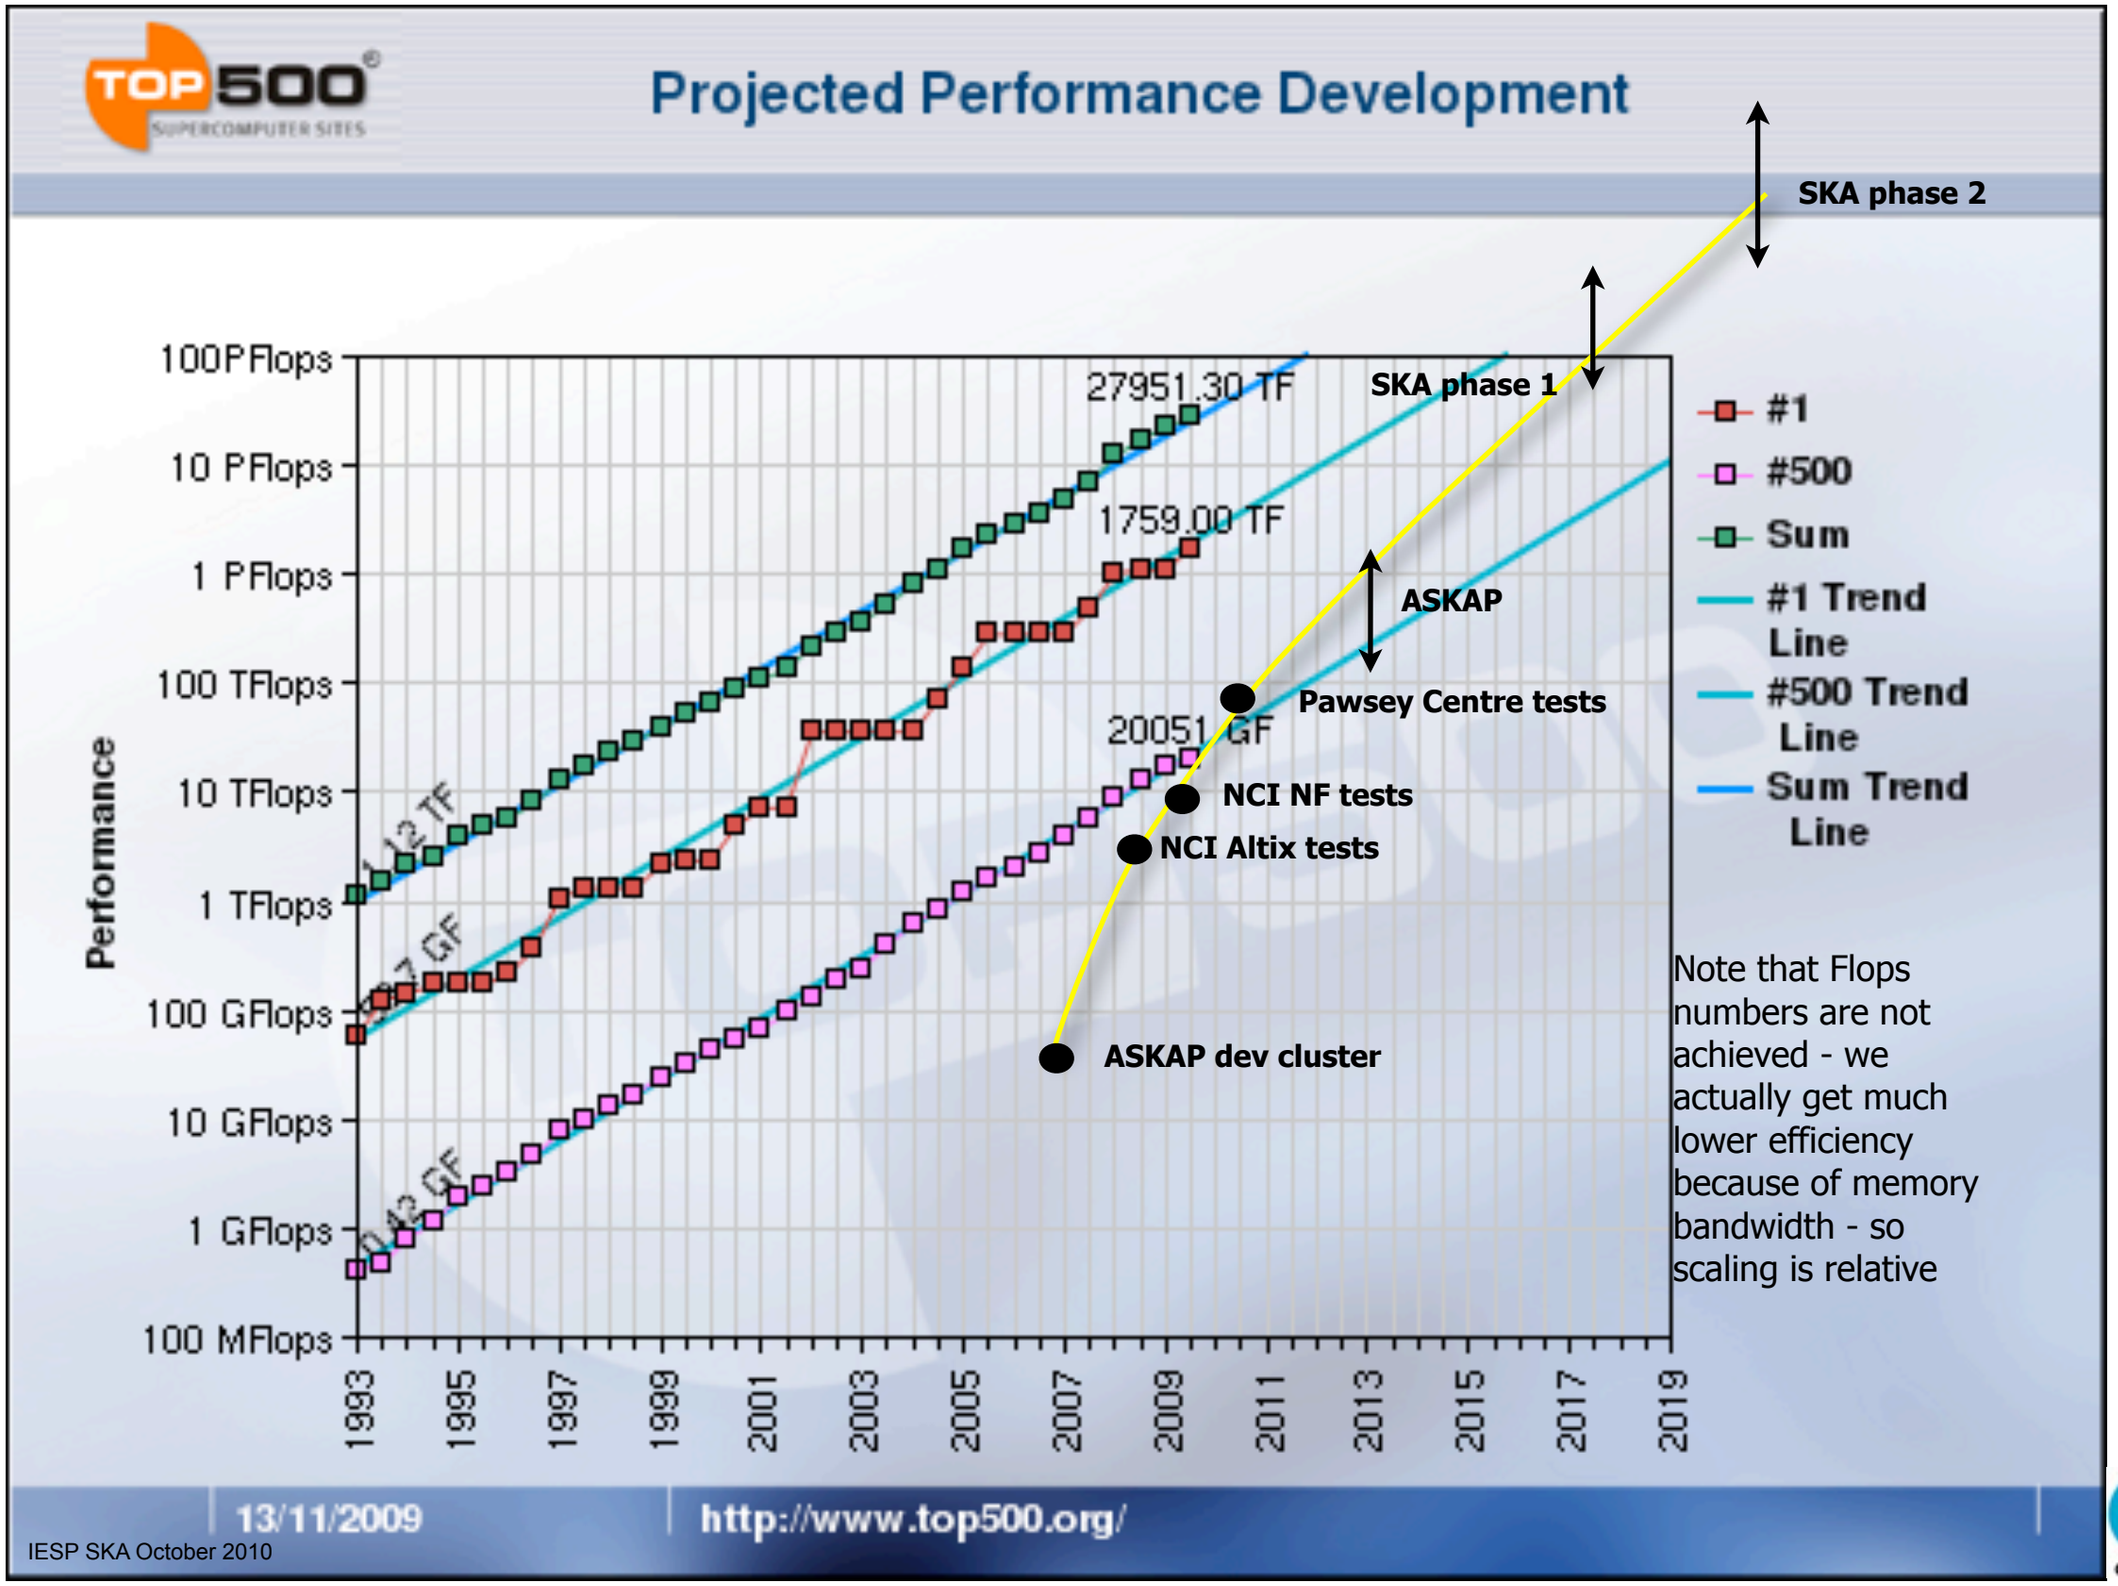

Note that Flops numbers are not achieved - we actually get much lower efficiency because of memory bandwidth - so scaling is relative

Does SKA meet the criteria for CDV

Petascale or near-petascale application today with a demonstrated need for exascale performance	•~ 100TF by end 2010 •~ 100TF, 100PB/year by 2012 •~ 1PF by 2014
Significant scientific goals in an area that is expected to be a scientific or societal driver for exascale computing	•Astronomy is widely recognised as a societal good and technology driver •Demonstrated by successful funding
Realistic and a definable set of steps to exascale that can be mapped out over 10 years or less	•Pathfinders (e.g. ASKAP, LOFAR) •SKA1 •SKA2
Community experienced in algorithm, software and/or hardware developments and is willing to engage in the exascale co-design process.	•Community is experienced •Exascale is necessary for science goals •Strongly motivated to participate in co-design
Modular and open enough to stimulate the development of additional modules addressing related questions in the area	•Astronomy is very open, nearly all OSS •Design, software shared across community
Fills a slot in the portfolio of extreme scale application needed to test all these dimensions	•Algorithms (near EP and non EP) •Data graphs, flow, and management

Conclusion

- Next steps
  - SKA now in initial costing phase, ending late 2011
  - Enter design and development phase in 2012
  - Construction of SKA1 starts 2016: 350M€ globally
  - Construction of SKA2 starts ~ 2020: 2G€ globally

- SKA yearly meeting next week in Oxford
  - I will report on IESP

- Why is SKA an important co-design vehicle?
  - Data-centric
  - Data management
  - Substantial opportunities for tailoring algorithms to hardware
  - Multiple opportunities for prototyping on real systems
  - One of the largest, most visible scientific instruments in this decade
  - Timescale matched to IESP

ATNF/ASKAP
Tim Cornwell
ASKAP Computing Project Lead

Phone: +61 2 9372 4261
Email: tim.cornwell@csiro.au
Web: www.atnf.csiro.au

www.csiro.au

Thank you

Contact Us
Phone: 1300 363 400 or +61 3 9545 2176
Email: enquiries@csiro.au Web: www.csiro.au

CSIRO

Radio Astronomy: Square Kilometre Array

Tim Cornwell – CSIRO Astronomy and Space Science
tim.cornwell@csiro.au

Scientific and computational challenges	Software issues – 2012-2020
• Formation of the first stars; Cosmology and galaxy evolution; Tests of GR; Cosmic magnetism; Cradle of life • Imaging, source finding, correlation, beam forming • 2012 (ASKAP): 36 dishes, 6192 feeds • 2016-2020 (SKA1): ~ 250 dishes	• Management of very large data flows • Software/algorithm solutions to widening flops vs memory-bandwidth gap • Development of new algorithms or methods of processing data • Scaling existing techniques to 100,000,000+ cores is not viable •"One-touch processing"

Software issues – 2010-2012	Impact of last machine change (10's of Gflops to 100 Tflops)
• Adapting software stack aimed at batch/ simulation type jobs to real-time data processing • Low computational intensity of algorithms, limited by memory-bandwidth • High input data rates • Managing large datasets (parallel-IO, archiving, transfer)	• Existing algorithms still useful, only adaptations necessary (1GB memory/core) • Required multi-disciplinary teams (HPC + science domain expertise) • Problem is mostly embarrassingly parallel to (O)10,000 cores • Optimizing I/O patterns critical

Demonstration of full range of SKA imaging

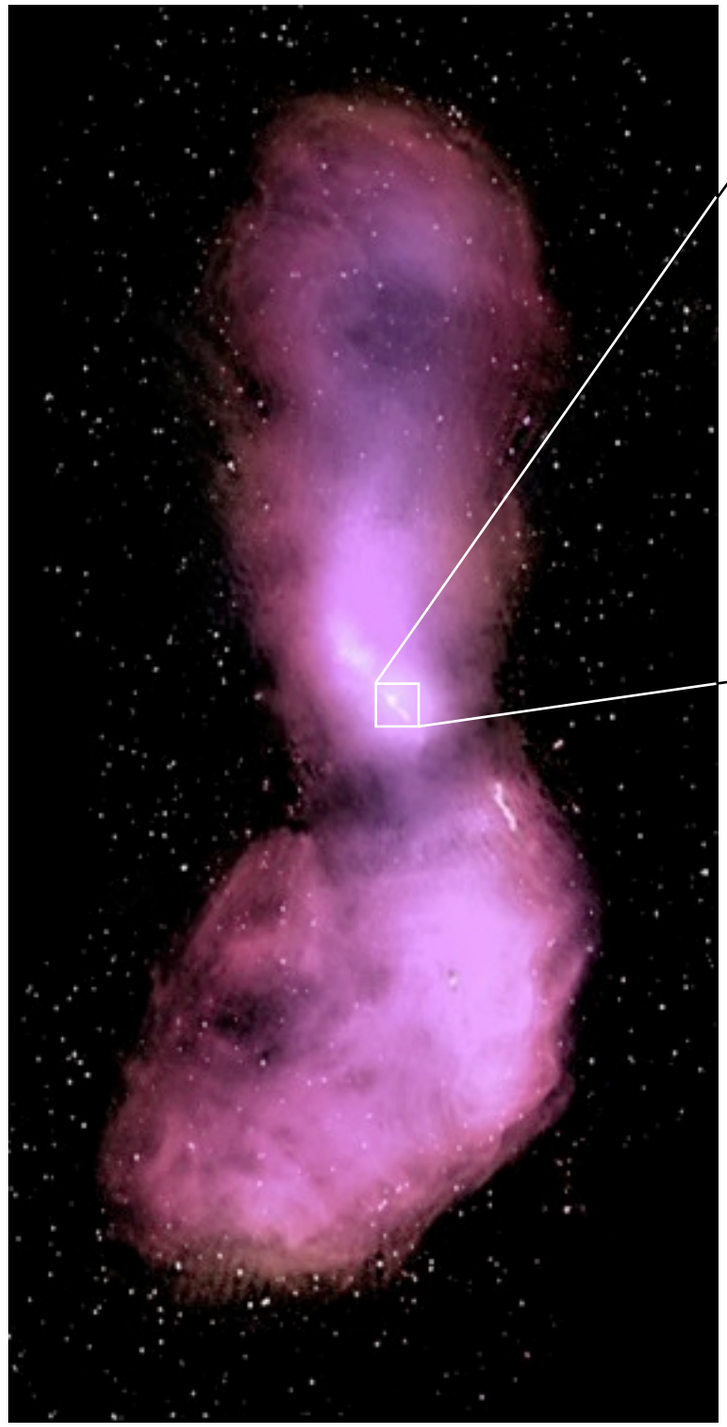

600 kpc

Feain, Cornwell, Ekers

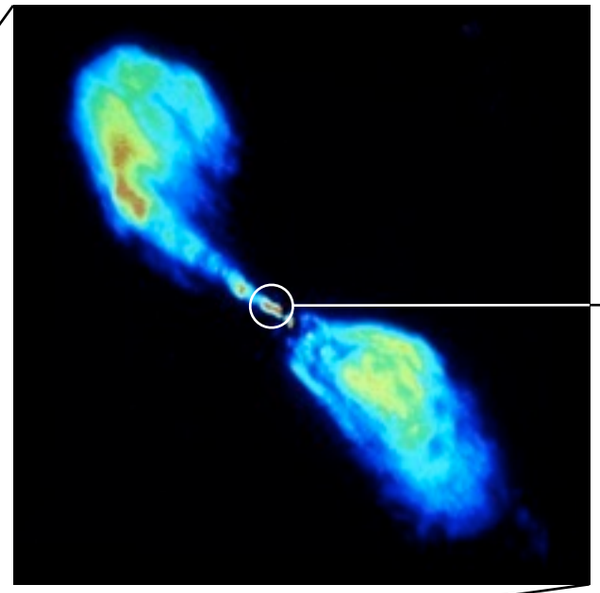

10 kpc

Burns & Clark

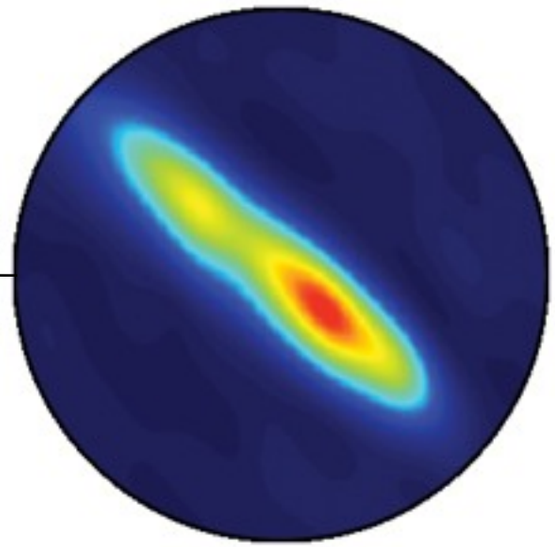

1 pc

Tingay et al.

LBA + ASKAP + Warkworth observations of nuclear region of Centaurus A at 21cm.