Big Data and Extreme Computing for the Square Kilometre Array

Tim Cornwell, Square Kilometre Array
EC: Ben Humphreys, CSIRO
BD: Andreas Wicenec, ICRAR

The Square Kilometre Array

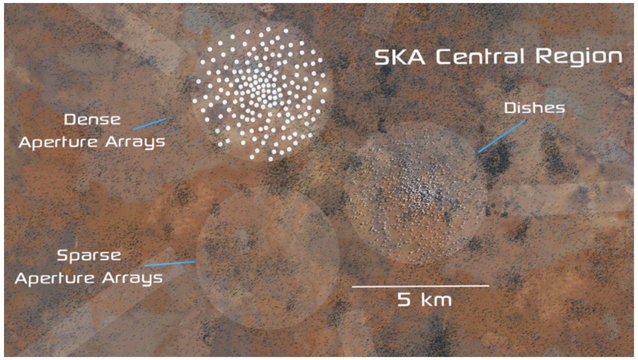

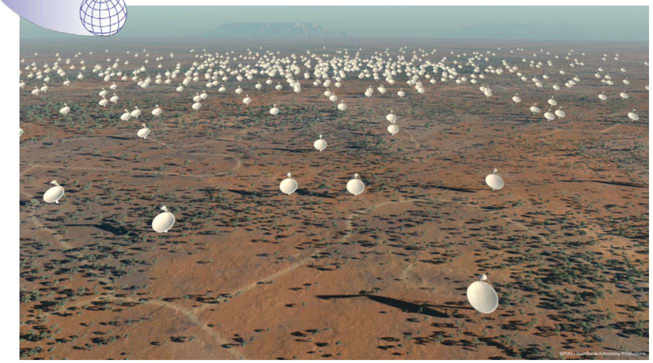

- 2020 era radio telescope
- Very large collecting area (km^2)
- Very large field of view
- Wide frequency range (70MHz - 25 GHz)
- Large physical extent (3000+ km)

- International project
- Telescope sited in Australia and South Africa
- Headquarters at Jodrell Bank, UK
- Multiple pathfinders and precursors now being built around the world
- Now entering pre-construction phase

Exploring the Universe with the world's largest radio telescope

Phase 1: 2016 – 2020: €400M (2007)

SKA1_Mid + MeerKAT

SKA1_Low

SKA1_Survey + ASKAP

	SKA Element	Location
Dish Array	SKA1_Mid : 190 x 15m dishes + SPFs	RSA
Low Frequency Aperture Array	SKA1_Low : 280 Aperture array stations	ANZ
Survey Instrument	SKA1_Survey : 60 x 15m dishes + PAFs	ANZ

Phase 2: 2020 – 2025: €1.2B (2007)

SKA2 Mid Dish

SKA2 Mid AA

SKA2 Low

	SKA Element	Location
Low Frequency Aperture Array	SKA2_Low	ANZ
Mid Frequency Dish Array	SKA2_Mid	RSA
Mid Frequency Aperture Array	SKA2_Mid_AA	RSA

SKA1 data products and transformations

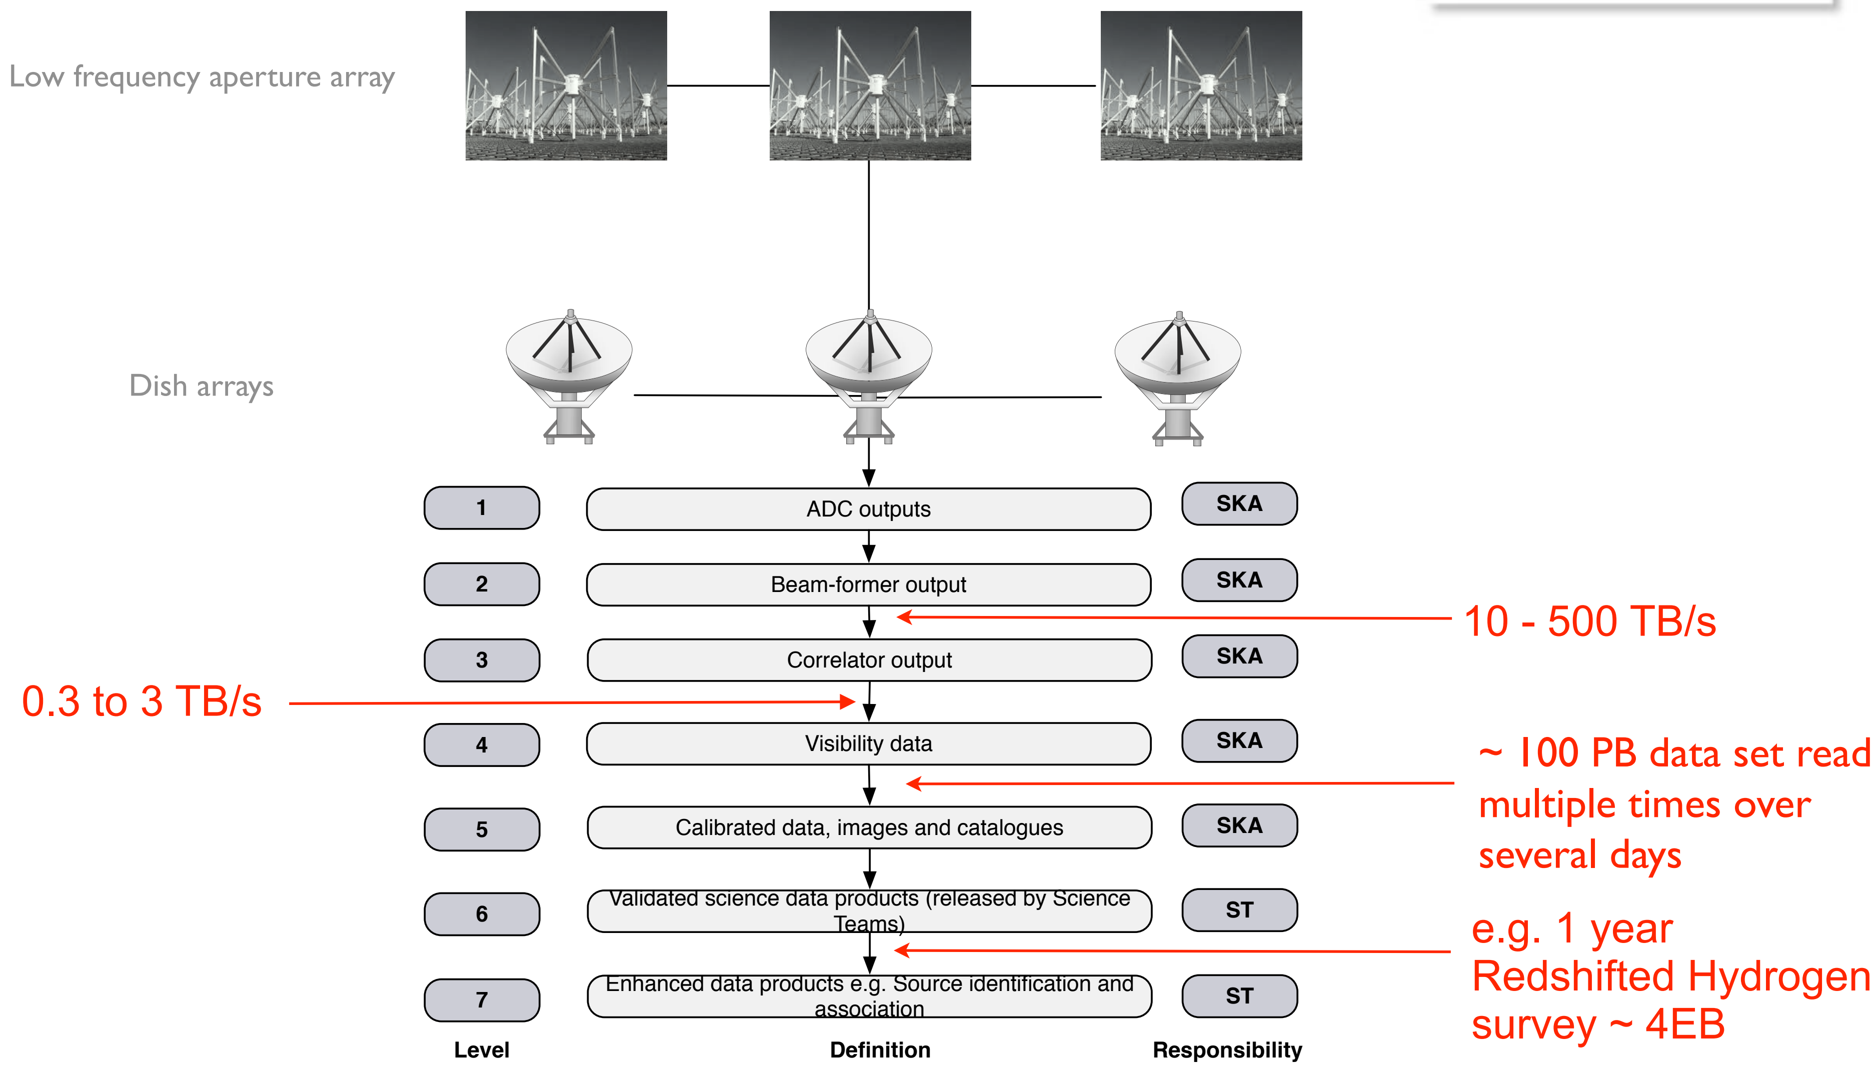

Example of cost model

Baseline	Ingest	Calibration	Continuum	Spectral_Line	Transient	All
10000	€213,303	€4,030	€73,835	€105,230	€4,030	€400,427
20000	€426,607	€28,955	€123,701	€463,022	€28,955	€1,071,239
30000	€639,910	€93,987	€231,008	€1,123,718	€93,987	€2,182,610
40000	€853,213	€218,314	€415,397	€2,121,029	€218,314	€3,826,267
50000	€1,066,516	€421,021	€696,314	€3,481,560	€421,021	€6,086,433
60000	€1,279,820	€721,416	€1,093,516	€5,227,707	€721,416	€9,043,876
70000	€1,493,123	€1,138,602	€1,626,417	€7,379,014	€1,138,602	€12,775,757
80000	€1,706,426	€1,691,733	€2,314,500	€9,952,941	€1,691,733	€17,357,332
90000	€1,919,730	€2,399,961	€3,177,228	€12,965,347	€2,399,961	€22,862,227
100000	€2,133,033	€3,282,435	€4,234,050	€16,430,813	€3,282,435	€29,362,766

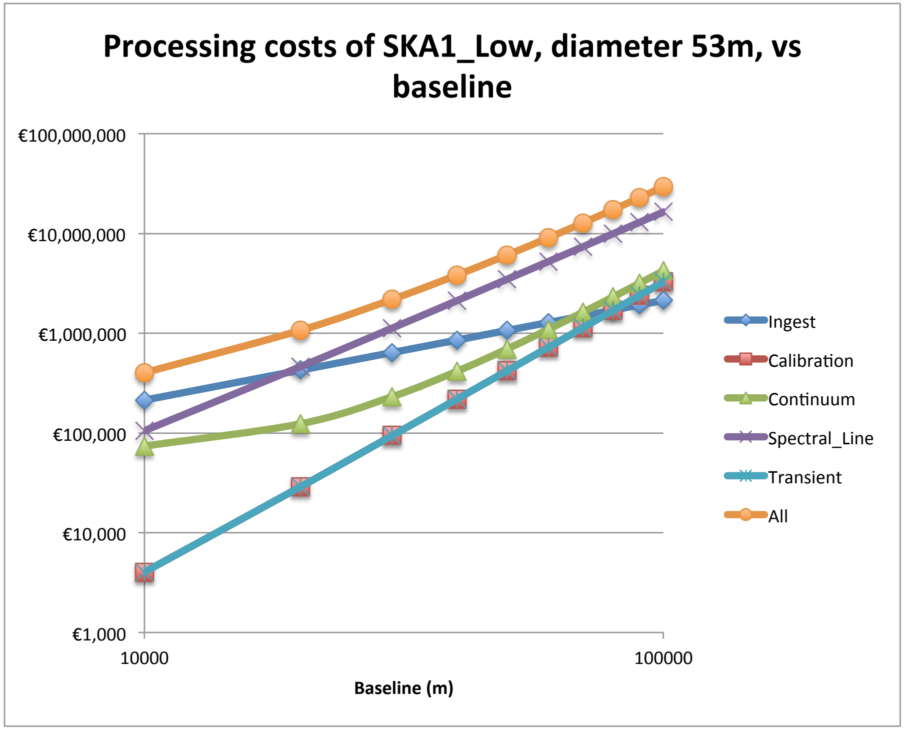

- 2019 costs
- Based on four kernel model of processing
- Expect to update this model continuously over preconstruction
- Poor on cost of data movement

Architecture

- What architectural changes are needed?
  - Tighter integration of storage, computing and networking.
    - The classical split between storage and computing requires that large amounts of data have to be moved even for very little computing.
  - More intelligent, hierarchical object storage and application driven networking
    - Could potentially perform in-storage filtering and/or on-stream transformations.
  - Dynamic integration of several memory and cache levels into intelligent data movement/pre-fetch agents.
  - Cassandra and Hadoop work well for some use cases

Workflows

- Describe a forwarding-looking workflow
  - Our workflows are mainly data reduction
    - Ingest, editing, calibration, imaging, source finding, analysis, archiving
    - With some iteration
    - Run constantly (as the telescope observes)
- What software is missing to support your workflow?
  - Tighter integration between data movement services and compute scheduling
  - Observability tools (monitoring) for data flow within systems
  - Something like Infosphere streams, Twitter Storm,...

Taxonomy

- Outline how you use your data
  - BDEC system is part of the telescope
  - Telescope becomes adaptive to cancel calibration effects
  - Steps are: acquire, edit, calibrate, transform, analyse, with iterative cycles
  - Too much data to allow guiding by humans
  - But analysis step requires some human guidance and performance
  - Analysis rich in visualization, feature identification, catalog queries
  - Survey science

Taxonomy

- Data-driven mini-application?
  - Not yet but will be developed
- Cross cutting concerns
  - Advanced I/O optimised data formats (e.g. ADIOS: http://www.olcf.ornl.gov/center-projects/adios/)
  - Usability of systems for astronomers. Barrier to entry is getting ever higher. Astronomers spending more time wrestling with systems and less doing core research.
  - Turnaround time, the ability to do quick iterations on petascale datasets just like we do on gigascale datasets.
    - Fast iteration == productivity

Taxonomy

- Data flow
  - Waterfall with successive phases of refinement and lossy compression
  - Processes simple at beginning of flow, less so at end
  - Some phases (calibration and imaging) required multiple (10?) passes through data
  - Analysis requires long residency of result in “working set” : days to weeks
  - Arithmetical complexity low throughout

Software

- What software are you currently using to manage and explore your data?
  - Software is largely bespoke
    - CASA, CFITSIO, WCSLIB
    - ASKAPSoft
    - SKA expects to spend about 150 FTE-years
- What algorithms and software libraries/tools need development and improvement to address your big data needs?
  - Synthesis calibration and imaging algorithms must scale up ~ 10,000
- Problems without a full solution
  - Fault tolerance

Interoperability challenges

- Provenance, etc.
  - Well in hand from IVOA work
- Semantics
  - There is a semantics working group in the IVOA and they have produced a complete ontology of astronomical object types
    - http://wiki.ivoa.net/internal/IVOA/IvoaSemantics/WD_2007-02-19.pdf
  - All of the IVOA standards use the UCD standard vocabulary.
    - A UCD does not define the units nor the name of a quantity, but rather "what sort of quantity is this?"
  - Implementation slow
  - Need a proper, controlled mechanism of machine-readable semantic tagging of quantities and data objects