Continuum surveys and computational issues

Tim Cornwell, Square Kilometre Array

How SKA processes data

- Science Data Processing system is <u>part</u> of the telescope
  - Only one system per telescope
- Data flow so large that dedicated facility is needed
- Telescope becomes adaptive to e.g. 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 <u>and</u> performance
- Analysis rich in visualization, feature identification, catalog queries
- Survey science

SKA data processing rates

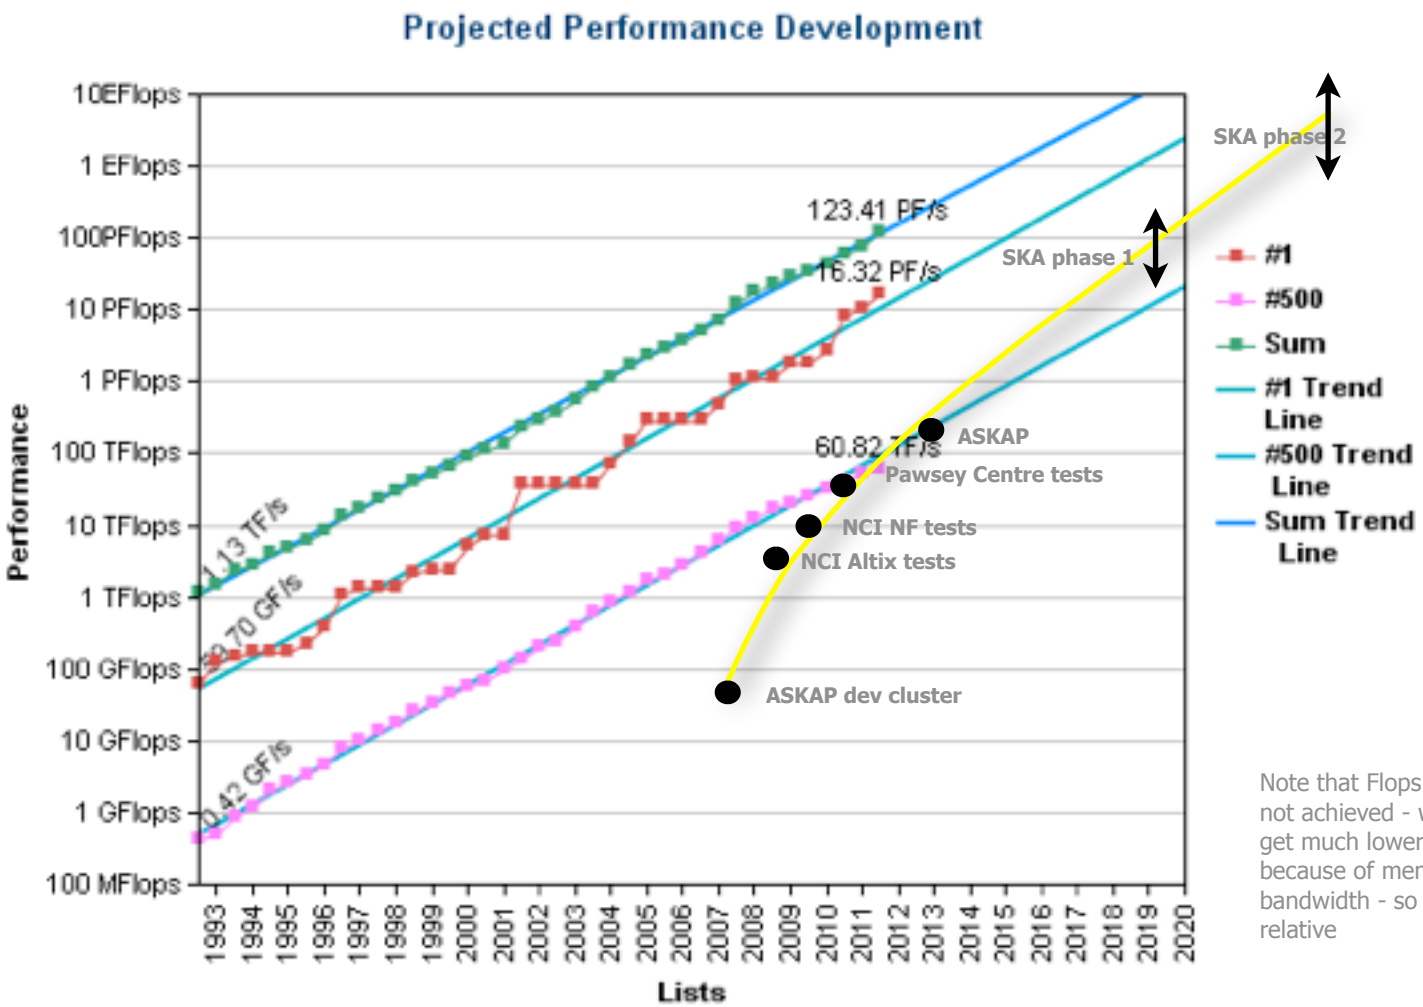

Note that Flops numbers are not achieved - we actually get much lower efficiency because of memory bandwidth - so scaling is relative

SKA1 data products and transformations

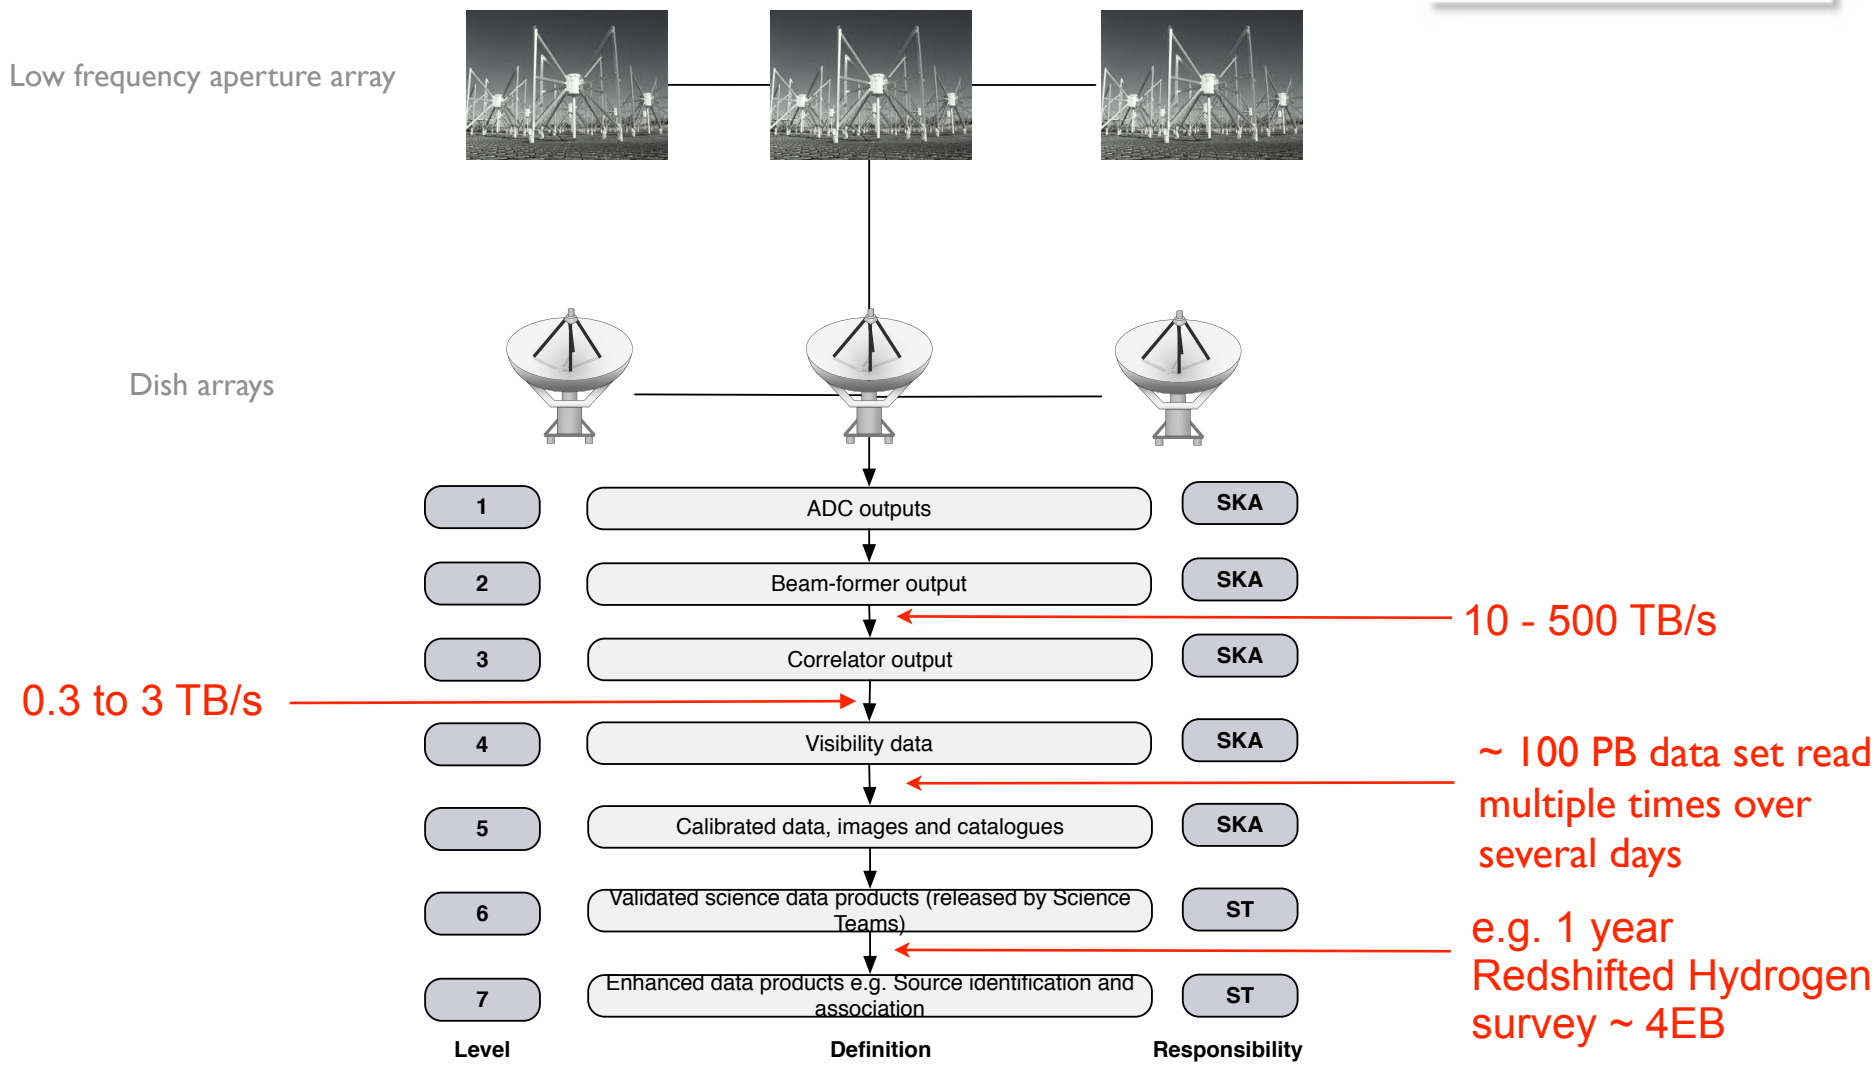

Exploring the Universe with the world's largest radio telescope

SKA data processing challenges

- Wide field imaging
  - ASKAP, MWA, LOFAR
- Imaging with aperture arrays
  - LOFAR, MWA
- Imaging with phased array feeds
  - ASKAP
- Imaging with wide bandwidth
  - JVLA , ASKAP
- Calibration and correction of direction dependent effects
  - LOFAR, MWA, , VLA, MeerKAT, ASKAP
- All organizations represented at SKA CALIM meetings
  - Challenges being addressed across community
- High performance computing lags

Costing calibration and imaging

- Don’t try to count flops, IO
- Scale existing software to SKA
- Use small benchmark programs to get info for scaling
  - Gridding, cleaning, reprojection, FFT
- Scale known configuration (and cost)

Calibration and imaging cost model

- Five pipelines
  - Ingest, Calibration, Continuum, Spectral line, Transients
- Steps in processing
  - Gridding and degridding visibility data
  - Multi-Scale Multi-Frequency CLEAN

$$T_{clean} = \mu_{clean} N_{scales} N_{Taylor} N_{iterations} N_{psf}^2$$

- Resources
  - Processing
  - Memory
  - Fast storage

Gridding/degridding model

$$T_{ws} = \mu_{grid} N_{vis} \left(\rho^2 \left(\frac{w_{rms}}{w_{max}} \right)^2 R_F^2 + R_A^2 \right) + \mu_{FFT} \left(2\rho R_F \left(\frac{T_{obs}}{T_A} \right) + N_{int} \right) N_{pixels}^2 \log_2 \left(N_{pixels}^2 \right)$$

$$+ \mu_{reproj} \frac{N_{pixels}^2 h_{obs}}{\rho}$$

- AW snapshots algorithm
  - Identify best-fit plane in uvw space
  - Use AW Projection to move points onto this plane
  - Update plane at rate chosen to minimize total work
  - Some projection mandated by Earth curvature
  - Correct for coordinate distortion of each image plane
- Superior to snapshot imaging and AW Projection
  - Less CPU, memory
- Diagonal or general Mueller matrices
- Update or change model in the future as appropriate

CPU-based scaling numbers

- Performance measured by four numbers

$$\mu_{wp}, \mu_{FFT}, \mu_{reproj}, \mu_{clean}$$

- Can be benchmarked by small programs
  - tConvolve, tHogbomClean
- Numbers shown are from Pawsey Centre ASKAP Real Time Computer
  - 200TF system costing €3.2M
- Expect SDP to update cost model regularly during pre-construction and construction
  - Track vendors
  - Track scaling work

μ_{grid}	8.0e-13 s
μ_{FFT}	1.1e-11 s
μ_{reproj}	9e-10 s
μ_{clean}	1.0e-12s
Cost (2012)	€1.6M

GPU-based scaling numbers

- Somewhat speculative
- Grid, clean measured
- FFT, reproject scaled
- Substantially better than CPU
- Use these numbers

μ_{grid}	8.0e-14 s
μ_{FFT}	1.1e-12 s
μ_{reproj}	9e-11 s
μ_{clean}	2.0e-13s
Cost (2012)	€1.6M

Data handling

- Single pass processing is insufficient to meet science goals e.g. dynamic range
  - Algorithms are non-linear
- Data must be buffered for ~ days to allow multiple calibration passes
- Require large multi-day visibility data buffer for all telescopes
- Assume average throughput must = 100%

Typical computing platform costs

- Based on GPU scaling
- Runs all pipelines
- Mid should have smaller maximum baseline

Telescope	Diameter	Baseline	Processor	Memory	Storage	Total
	(m)	(km)	(M€)	(M€)	(M€)	(M€)
SKA1_AA_Low	35	100	14.3	0.9	1.4	16.6
SKA1_Mid	15	200	59.3	19.1	8.9	87.3
SKA1_Mid	15	100	19.2	4.8	1.3	25.3
SKA1_Survey	15	50	11.9	1.6	0.8	14.3

AA-Low vs baseline, GPU

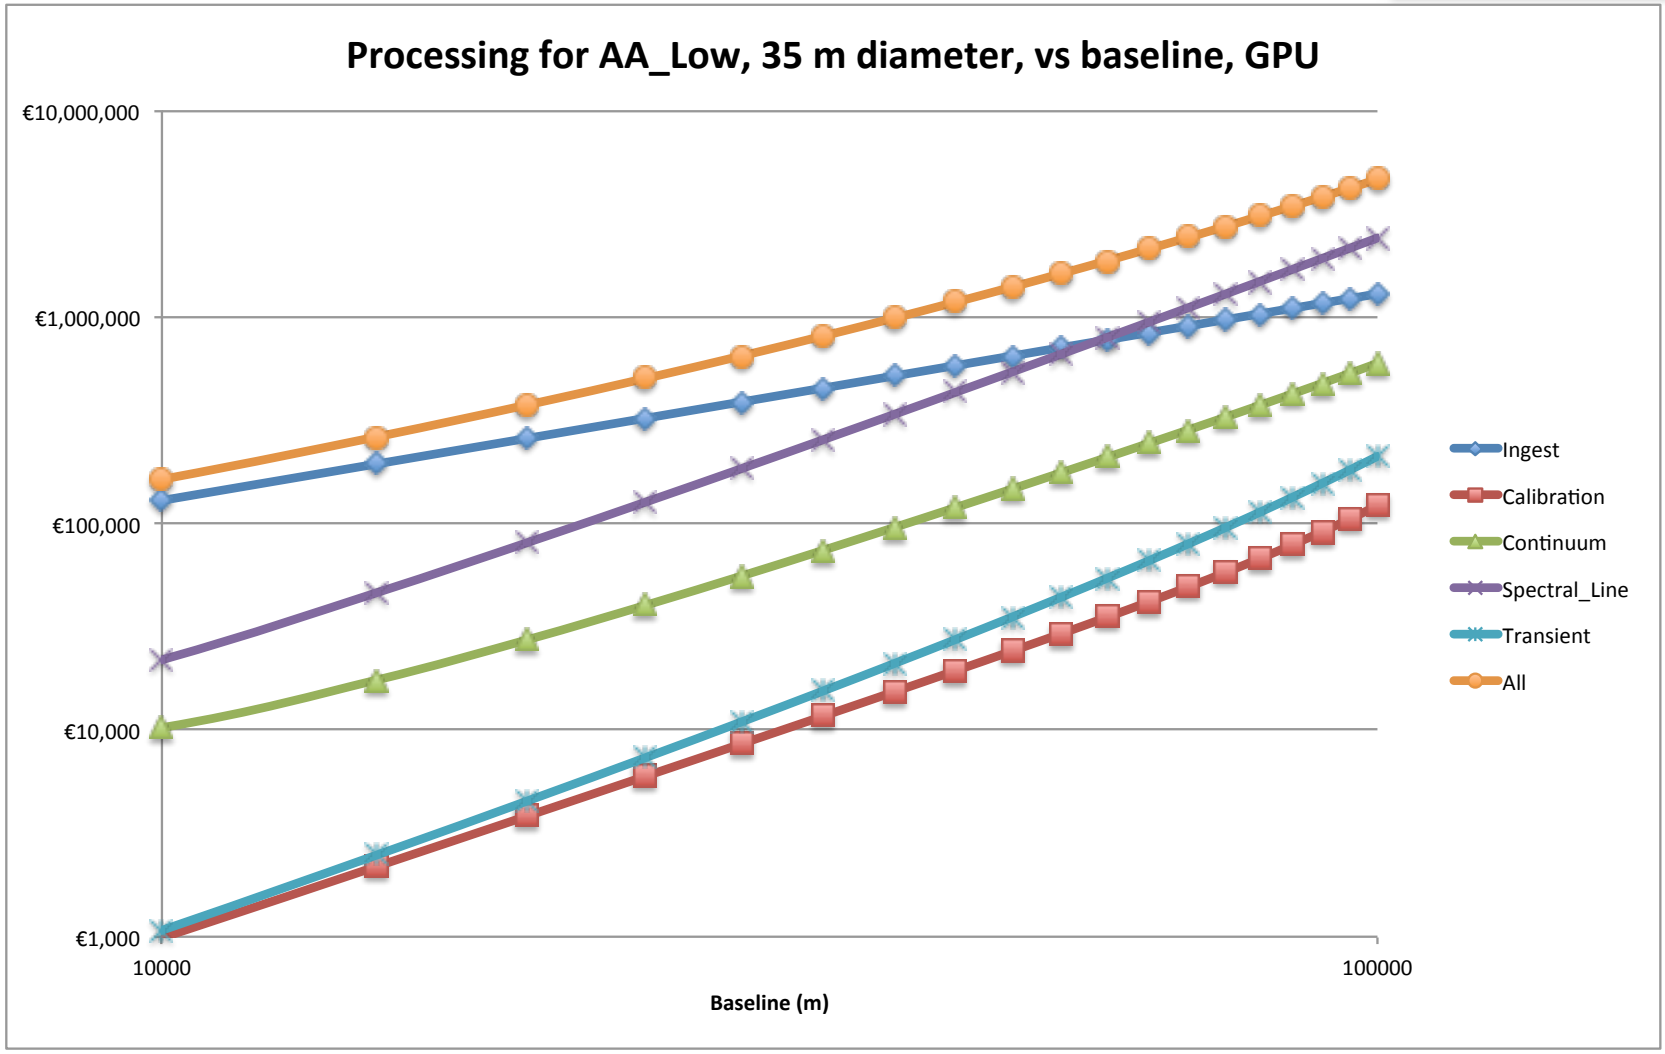

Why is this so good?

- New algorithm: w snapshots
  - Has better performance and scaling than w projection
- GPUs have better cost/performance than CPUs used in Pawsey
  - Intel MIC systems may be close
- Completely untested at the relevant scale

Caveats

- Assumes Moore's Law to 2019
- Assumes that processing can be scaled up factor of 1000 - 10000
- Assumes that we don't need radically new processing
- True for ASKAP from 2007 to 2013

Scaling ASKAP Processing to ~10,000 cores, SKA must go to ~10,000,000 cores

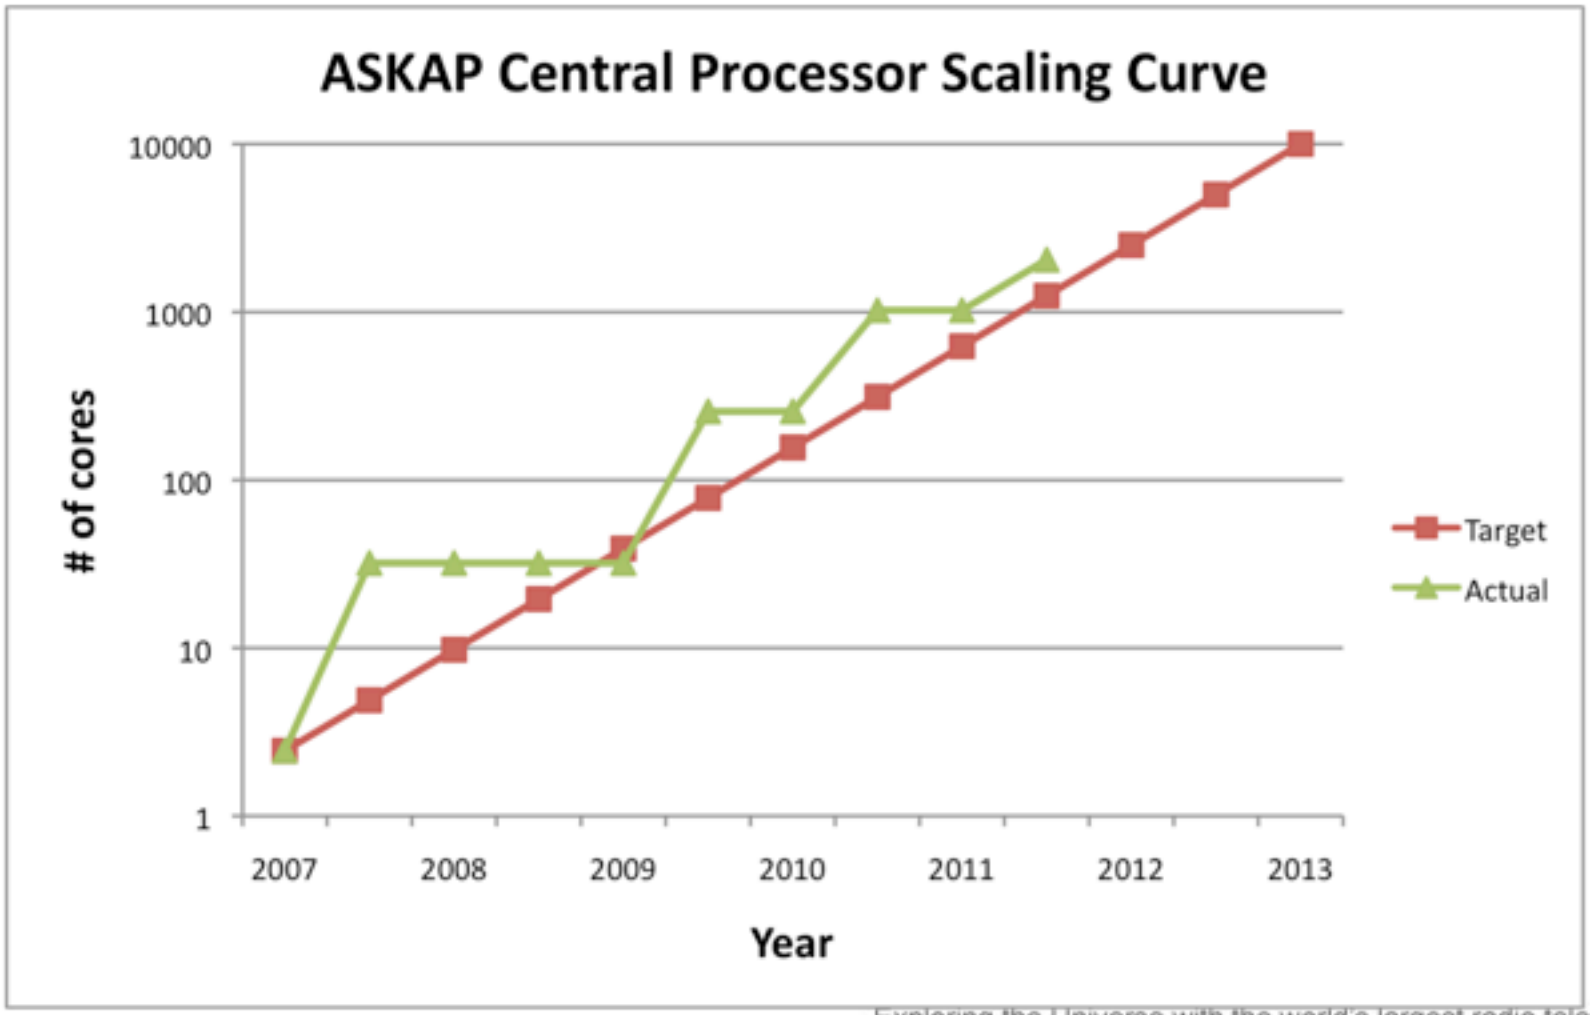

Exploring the Universe with the world's largest radio telescope

Processing capabilities

- Ignore scientific performance issues
  - e.g. noise level, dynamic range
- What can we do?
  - Fast surveys (several days)
  - Slow transients (correlator integration time 1s or longer)
  - Fast transients (1ms or more)

Fast continuum surveys (~ 3 days)

- SKA1-Mid
  - On-the-fly necessary
  - Slew up and down ~ 30 degrees in 6 minutes
  - No need to dump data faster:
    - 0.3s compared to 0.25s fixed pointing
  - Required SDP processing goes up a bit
    - Primary beam model must include slewing

Fast continuum surveys (~ 3 days)

- SKA1-Survey
  - On-the-fly necessary
  - Slew up and down ~ 30 degrees in 24 minutes
  - Required SDP processing goes up a bit
    - Primary beam model must include slewing

Slow continuum surveys (> 100 days)

- Mid and Survey
  - No obvious substantial hurdles compared to pointed observations
  - Storage requirements may be larger

Slow transient imaging

- Image entire field of view every correlator time time (~ 1 s)
- Image after having removing best model of sky
- Pass to specialised analysis software
- Cost can be calculated

$$T_{ws} = \mu_{grid} N_{vis} \left(\rho^2 \left(\frac{w_{rms}}{w_{max}} \right)^2 R_F^2 + R_A^2 \right) + \mu_{FFT} \left(2\rho R_F \left(\frac{T_{obs}}{T_A} \right) + N_{int} \right) N_{pixels}^2 \log_2 \left(N_{pixels}^2 \right)$$
$$+ \mu_{reproj} \frac{N_{pixels}^2 h_{obs}}{\rho}$$

- Generally acceptable cost
- ASKAP and VAST implementing this approach

Effect of antenna diameter on cost

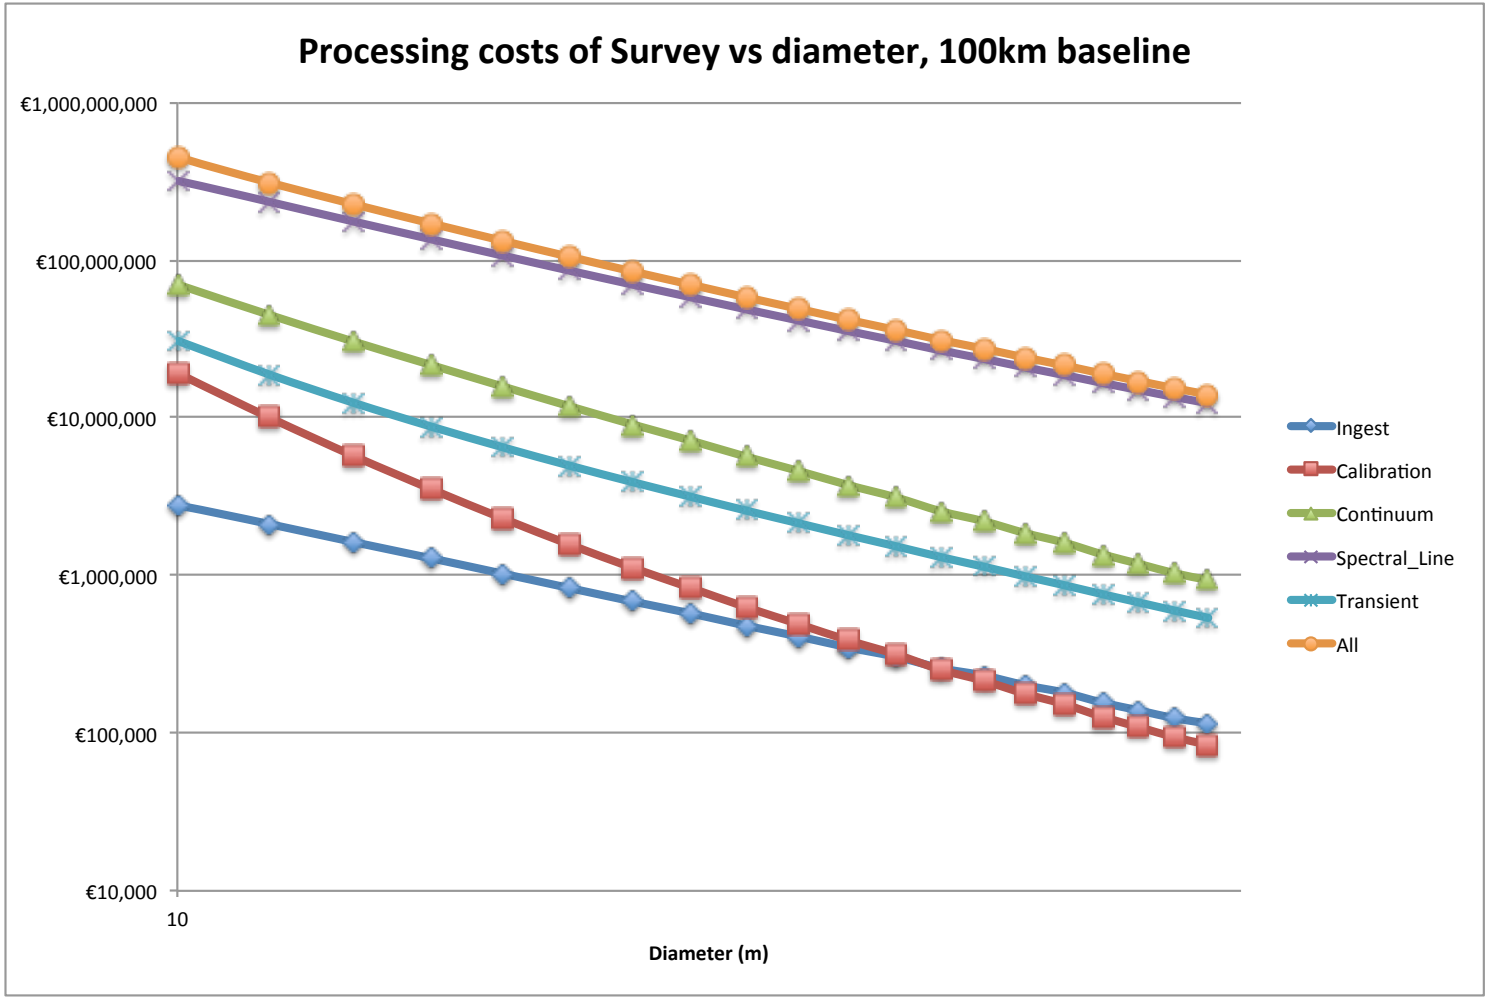

Exploring the Universe with the world's largest radio telescope

Effect of baseline on processing cost

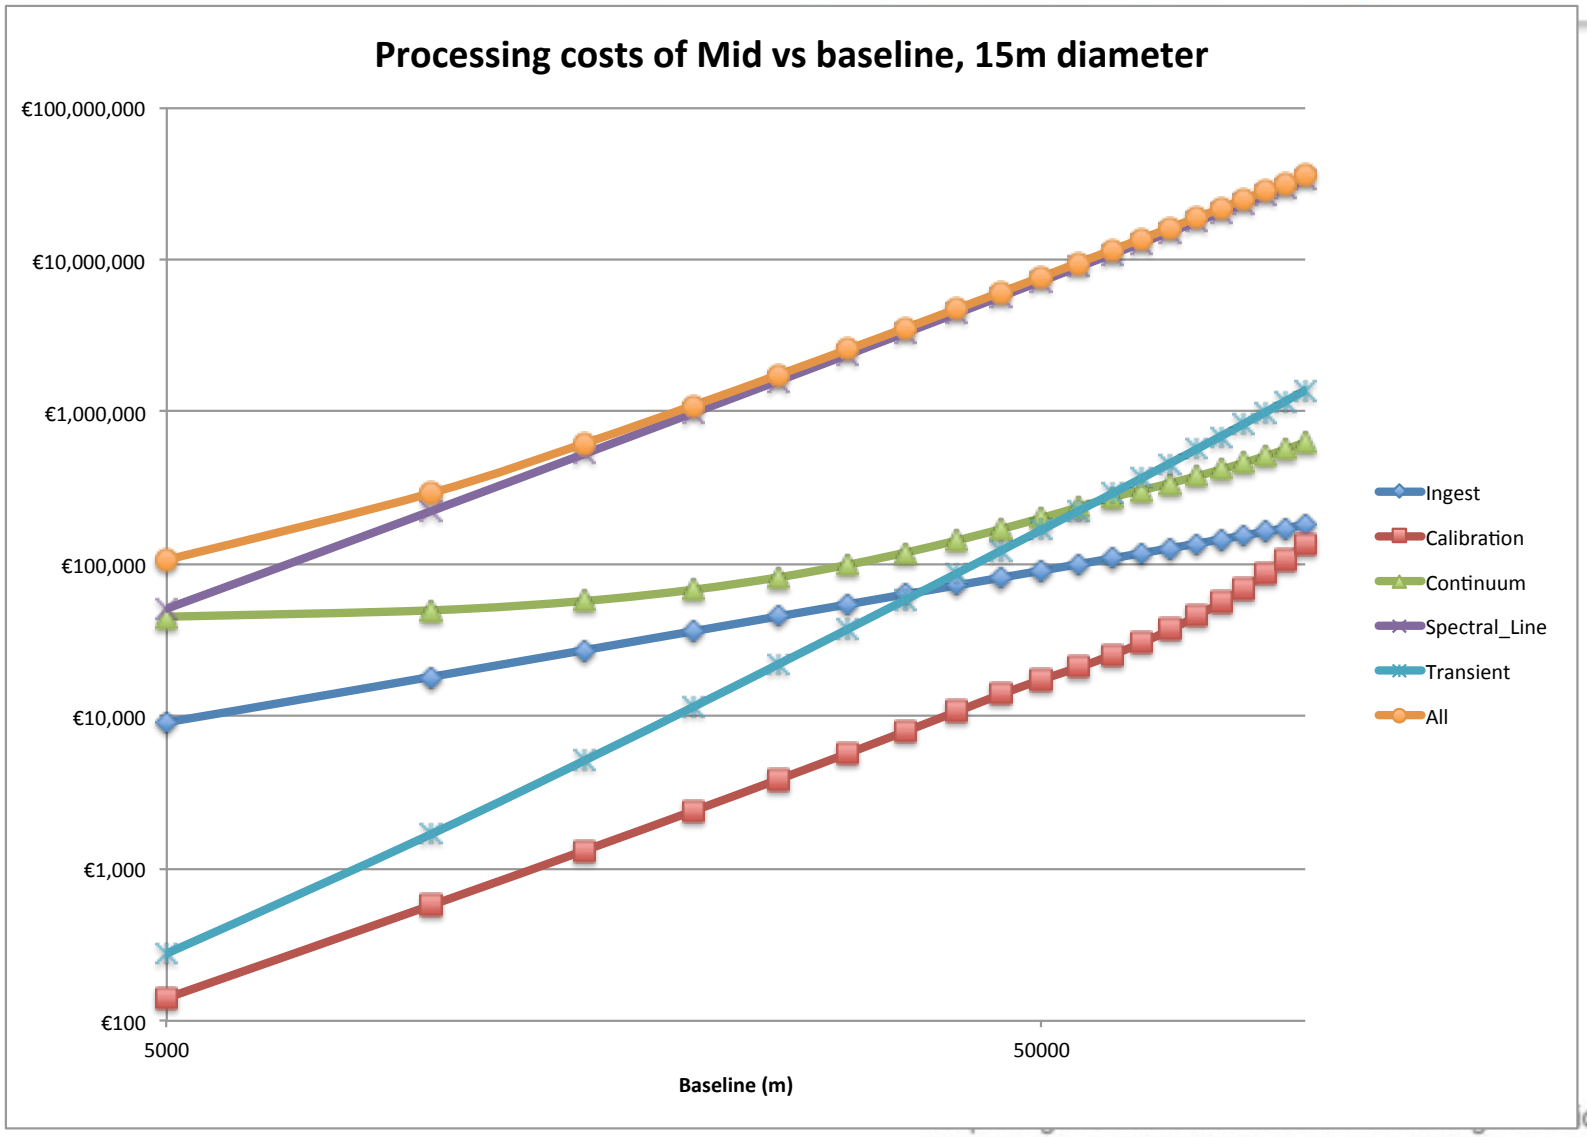

Processing cost for fast transients on SKA1-Mid

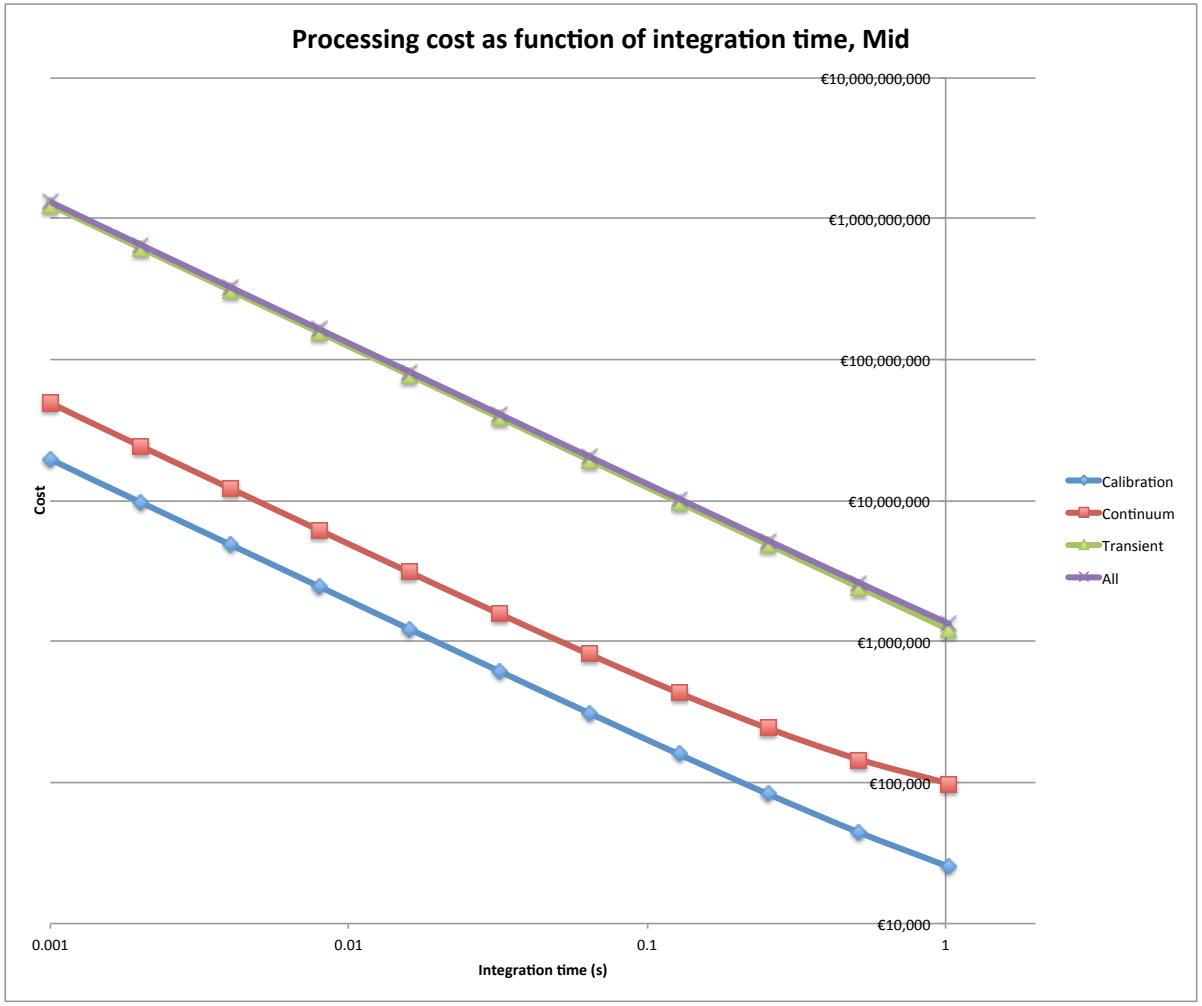

Exploring the Universe with the world's largest radio telescope

Accuracy of deconvolution

- Both for compact and extended structures
  - Shape of sources
  - Noise floor
- CLEAN has served us well!
  - Multiscale extended lifetime
- Plenty of evidence that Compressive Sampling does better
  - See e.g. papers by Wiaux and collaborators
- Need to see the effects on SKA size images
  - Scaling
  - Actual improvements on realistic fields
- Also need multi-frequency CS

Model for funding compute capacity

- Initial investment sufficient only for testing, commissioning, and early science
- Capacity will be incremented as required
- Part of Board strategy for cost cap

Summary

- Telescope computing capacity will grow as required
- Scaling work will probably continue into operations
- Very demanding hands-off processing model
- Surveys will be possible only after debugging a lot of difficult problems