SKA Data Flow and Processing – a key SKA design driver

Paul Alexander

SKA Overall System

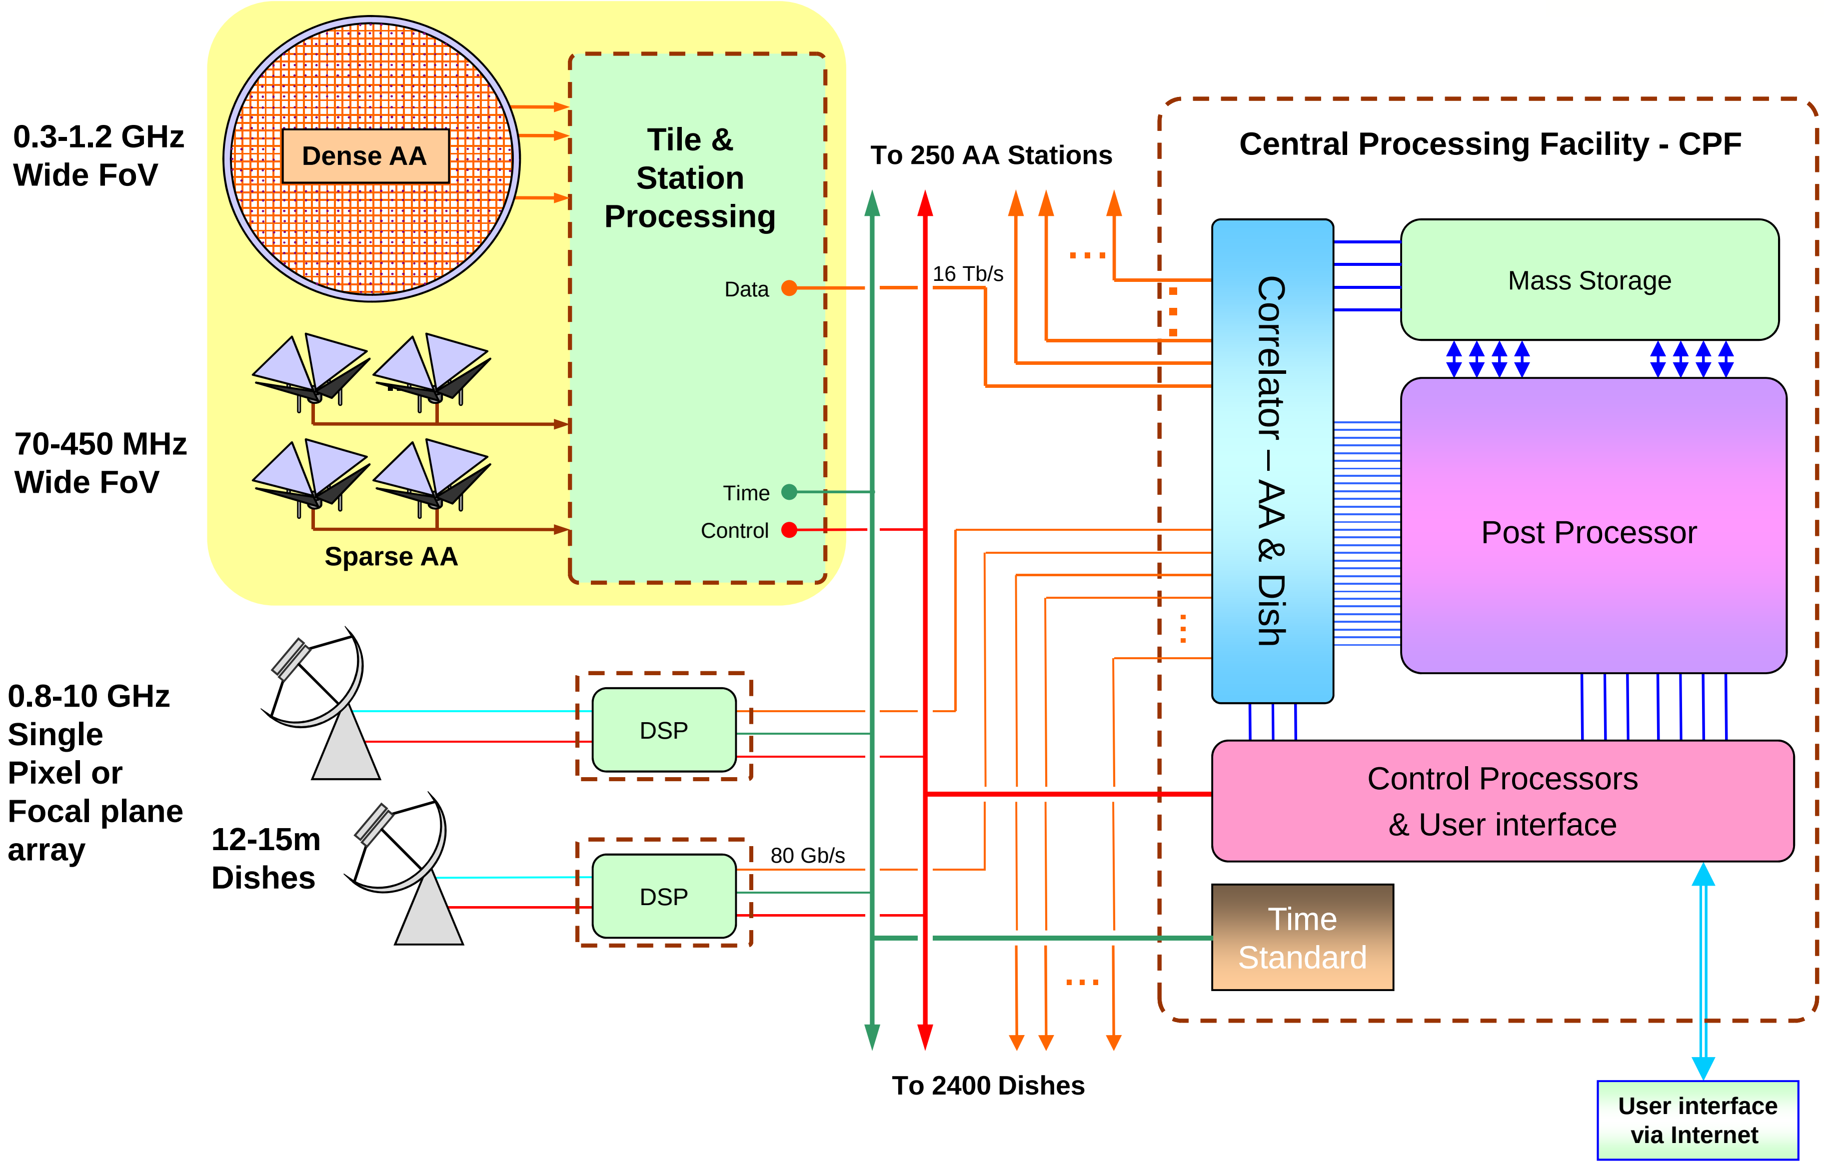

0.3-1.2 GHz
Wide FoV

Dense AA

70-450 MHz
Wide FoV

Sparse AA

Tile &
Station
Processing

Data

Time

Control

0.8-10 GHz
Single
Pixel or
Focal plane
array

12-15m
Dishes

DSP

DSP

80 Gb/s

To 250 AA Stations

16 Tb/s

To 2400 Dishes

Central Processing Facility - CPF

Correlator – AA & Dish

Mass Storage

Post Processor

Control Processors
& User interface

Time
Standard

User interface
via Internet

Some numbers to remember

Fastest HPC today – Roadrunner ~ 1.2 PFlop

Total internet data rate ~ 20 Tb/s

BUT on your desktop you can have ~1 TFlop

Data rates to the correlator

Data rate from each collector

$$G_1 = 2\,N_p\,\Delta f\,N_{bit}\,N_b = 4\,\Delta f\,N_{bit}\,N_b$$

$$N_b = \frac{1}{\Delta f}\int_{f_{max}-\Delta f}^{f_{max}} n_b(f)\,df$$

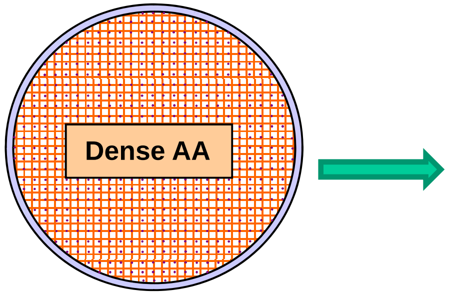

AA, Number of elements $N_e \sim 65000$; $N_b \ll N_e$ limited by data rate

250 sq-deg across band $N_b \sim 1200$ (Memo 100)

$G_1 \sim 16$ Tb/s

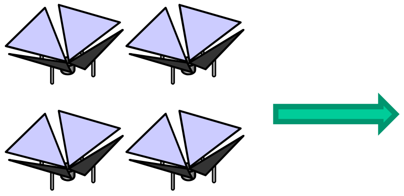

Sparse AA

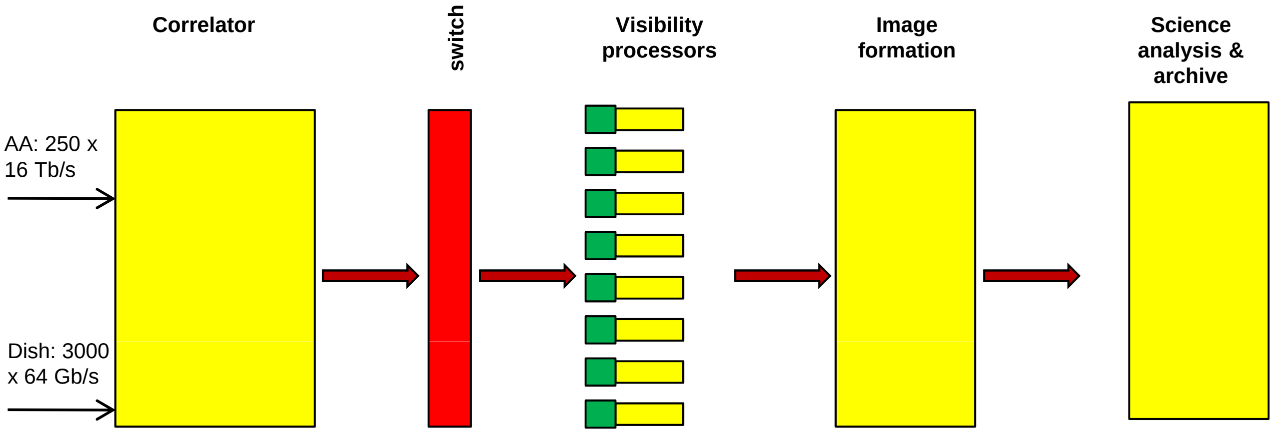

12-15m Dishes

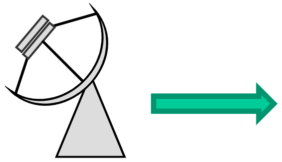

Dishes $G_1 \sim 64$ Gbs

PAFs FoV is constant across the band

$N_b \sim 7$ to give 20 sq-deg across the band $G_1 \sim 60$ Gb/s (Memo 100)

Data rates from the correlator

- Standard results for integration/dump time and channel width

$$\frac{\delta t}{s} = a_t \frac{D}{B} \sim 1200 \frac{D}{B} \qquad \frac{\delta f}{f} = a_f \frac{D}{B} \sim \frac{1}{10} \frac{D}{B}$$

- Naive data rate then given by

$$G = g(B) \frac{1}{2} N^2 N_p^2 N_b \frac{1}{\delta t} \frac{\Delta f}{\delta f} 2N_w \qquad G = g(B) N^2 N_w N_p^2 N_b \frac{1}{a_t a_f} \frac{\Delta f}{f} \left(\frac{B}{D}\right)^2$$

- Can reduce this using baseline-dependent integration times and channel widths

$$G = N^2 N_w N_p^2 N_b \frac{1}{a_t a_f} \frac{\Delta f}{f} \int_0^B n(b) \left(\frac{b}{D}\right)^2 \, db$$

$$= N^2 N_w N_p^2 N_b \frac{1}{a_t a_f} \frac{\Delta f}{f} \left(\frac{B}{D}\right)^2 \int_0^B n(b) \left(\frac{b}{B}\right)^2 \, db$$

Data rates from the correlator

- Standard results for integration/dump time and channel width

$$\frac{\delta t}{s} = a_t \frac{D}{B} \sim 1200 \frac{D}{B} \qquad \frac{\delta f}{f} = a_f \frac{D}{B} \sim \frac{1}{10}\frac{D}{B}$$

- Naive data rate then given by

$$G = g(B)\frac{1}{2}N^2 N_p^2 N_b \frac{1}{\delta t}\frac{\Delta f}{\delta f} 2N_w \qquad G = g(B) N^2 N_w N_p^2 N_b \frac{1}{a_t a_f}\frac{\Delta f}{f}\left(\frac{B}{D}\right)^2$$

- Can reduce this using baseline-dependent integration times and channel widths

$$G = N^2 N_w N_p^2 N_b \frac{1}{a_t a_f}\frac{\Delta f}{f} \int_0^B n(b)\left(\frac{b}{D}\right)^2 \, db$$

$$= N^2 N_w N_p^2 N_b \frac{1}{a_t a_f}\frac{\Delta f}{f} \left(\frac{B}{D}\right)^2 \int_0^B n(b)\left(\frac{b}{B}\right)^2 \, db$$

Data rates from the correlator

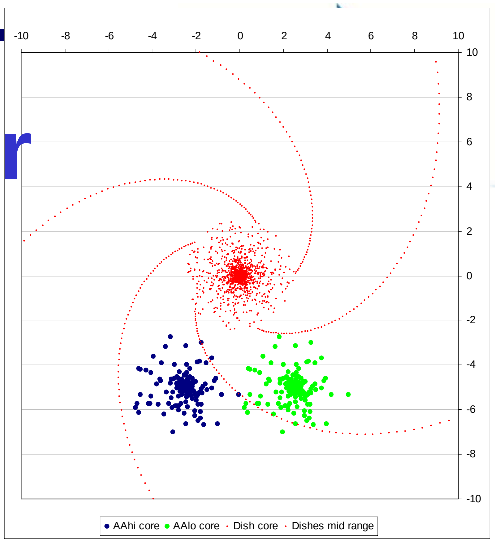

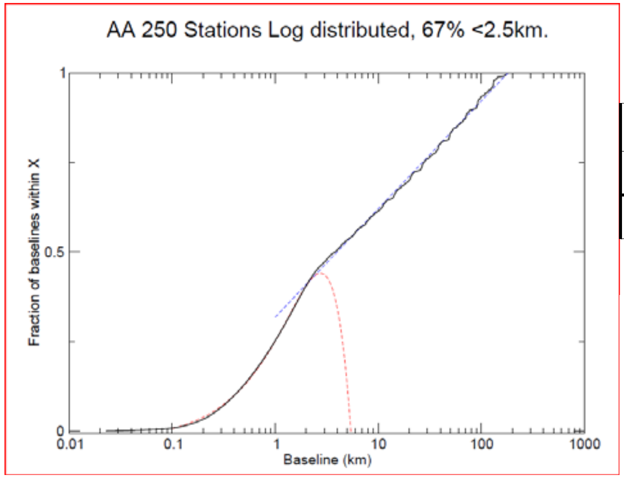

$$n(x) = \begin{cases} \beta / x & x \geq x_1 \\ \alpha & x < x_1 \end{cases}$$

	α	β	**B₁ / km**
Dishes	0.08	0.11	2.8
Aperture Array	0.33	0.13	2.2

$$g(x) = \int_0^x n(x')dx = \begin{cases} \beta \ln\left(\frac{x}{x_1}\right) + \alpha x_1 & x \geq x_1 \\ \alpha x & x < x_1 \end{cases}$$

Reduction in data rate using baseline-dependence (> 0.06)

1/3 $\quad B = B_1$

$$\frac{1}{2}\beta / \left(\beta \ln\left(\frac{B}{B_1}\right) + \alpha B_1\right) \qquad B >> B_1$$

Data rate out of correlator exceeds input data rate for 15-m dishes for baselines exceeding ~ 130km (36km if single integration time)

At best for dishes output data rate ~ input; AA’s reduction by ~10^4

Data rates from the correlator

Experiment				3000 Dishes + SPF		1630 Dishes + PAFS		250 AA stations	
Description	B_{max} (km)	Δf (MHz)	f_{max} (MHz)	Achieved FoV[1]	Data rate (Tb/s)	Achieved FoV[1]	Data rate (Tb/s)	Achieved FoV[1]	Data rate (Tb/s)
Survey: High surface brightness continuum	5	700	1400	0.78	0.055	15	0.11	108	0.03
Survey: Nearby HI high res. 32000 channels	5	700	1400	0.78	1.0	15	2.0	108	2.6
Survey: Medium spectral resolution; resolved imaging (8000)	30	700	1400	0.78	1.2	15	2.4	108	5.4
Survey: Medium resolution continuum	180	700	1400	0.78	33.1	15	66	108	14.1
Pointed: Medium resolution continuum deep observation	180	700	1400	0.78	33.1			0.78	0.15
High resolution with station beam forming[2]	1000	2000	8000	0.0015	33.4				
High resolution with station beam forming[3]	1000	2000	8000	0.0015	429				
Highest resolution for deep imaging[2]	3000	4000	10000	0.001	391				

Notes

1. Achieved FoV is at f_{max} and has units of degrees squared. For the AA and PAFs we calculate the data rate assuming it is constant across the band.
2. Assuming that for the dynamic range the FoV of the station only has to be imaged

3. Assuming that for the dynamic range the FoV of the dish must be imaged

Processing model

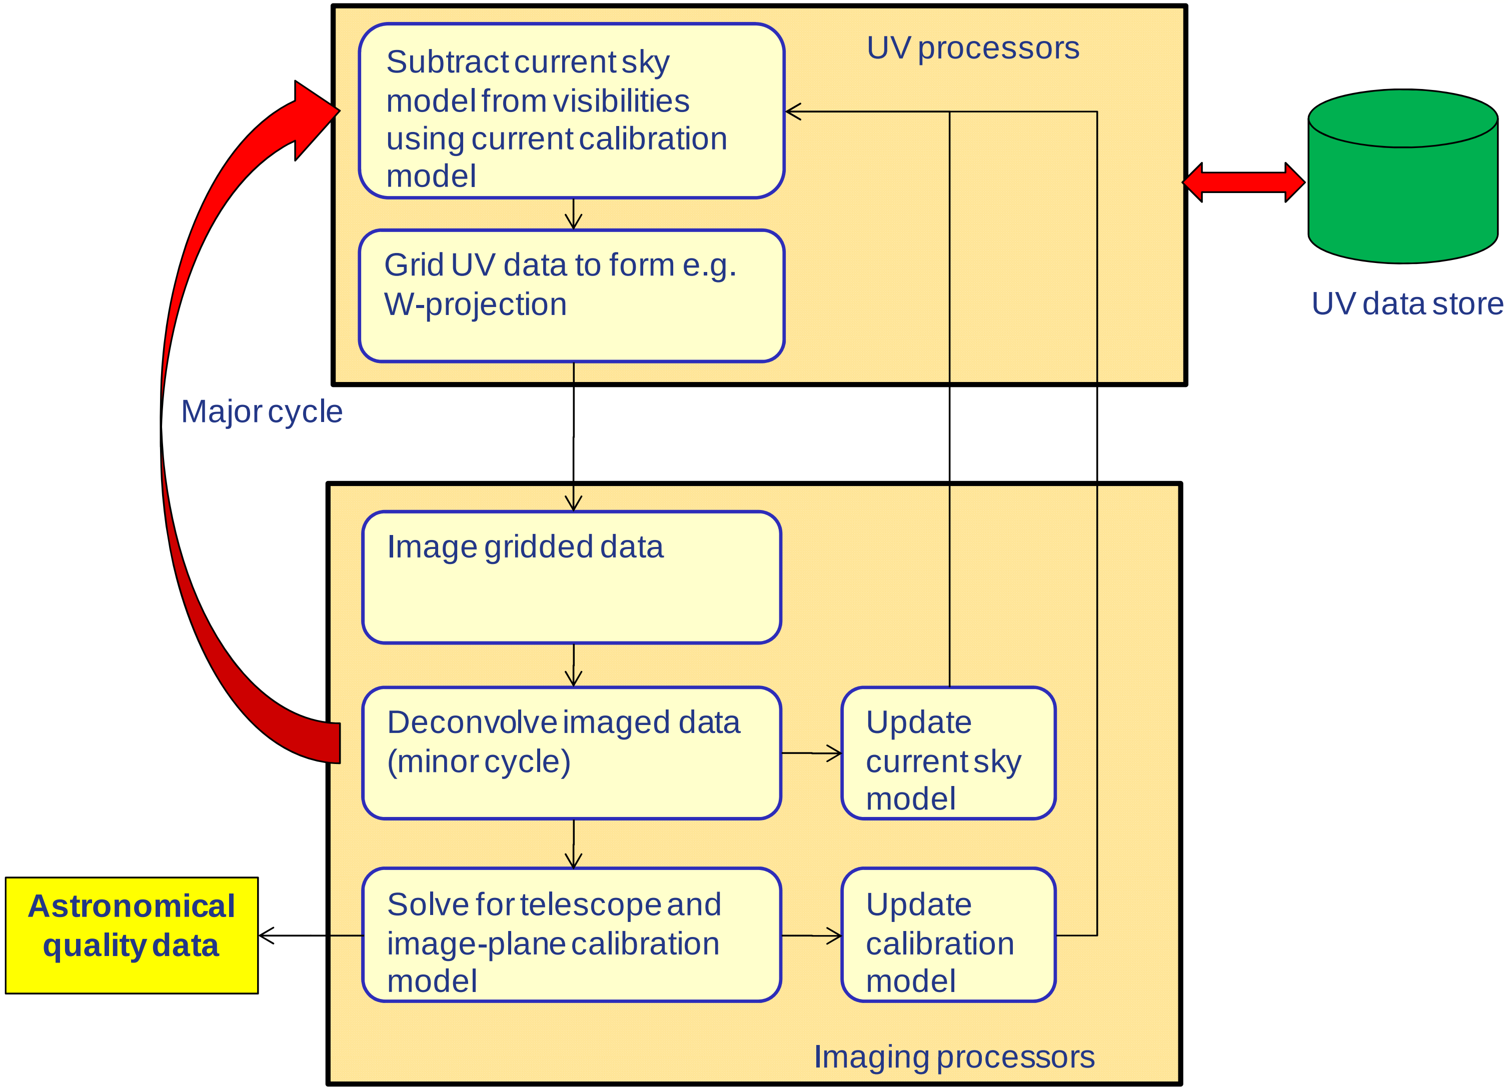

UV processors

Subtract current sky model from visibilities using current calibration model

Grid UV data to form e.g. W-projection

UV data store

Major cycle

Image gridded data

Deconvolve imaged data (minor cycle)

Update current sky model

Astronomical quality data

Solve for telescope and image-plane calibration model

Update calibration model

Imaging processors

Post correlator

- Processing on the raw UV data is highly parallel
- UV data needs to be buffered based on current algorithms
- UV processing includes:
  - Gridding and de-gridding from current sky model
  - Peeling and subtraction of current sky model
- Image size: $a^2 N_{ch} (B/D)^2 N_b$ Ratio UV to “image” data

$$\sim 0.06\, T_{obs}\, N^2 g(B) \frac{\Delta f}{f} \frac{1}{a_t a_f} \frac{1}{a^2} \frac{N_p^2}{N_{ch}} \quad \sim 210 \left(\frac{T_{obs}}{1\text{min}}\right) \left(\frac{N}{1000}\right)^2 \left(\frac{N_{ch}}{32000}\right)^{-1}$$

Major reduction in data rate occurs between UV data and image data

The SKA Processing Challenge

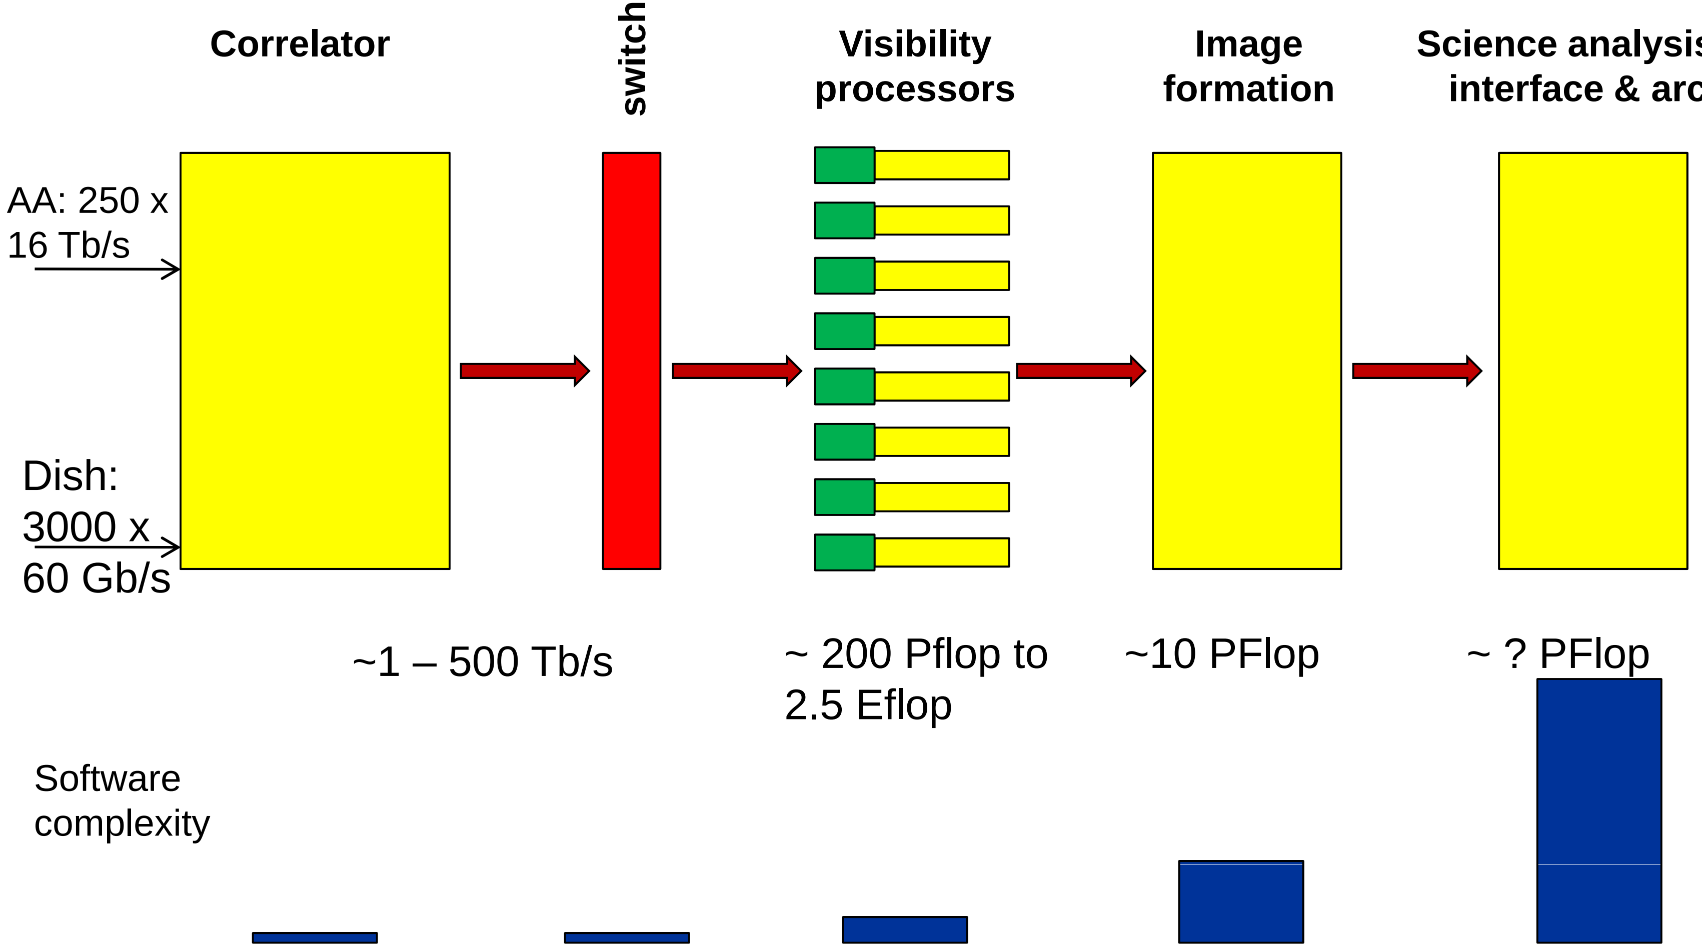

Correlator

switch

Visibility processors

Image formation

Science analysis, user interface & archive

AA: 250 x 16 Tb/s

Dish: 3000 x 60 Gb/s

~1 – 500 Tb/s

~ 200 Pflop to 2.5 Eflop

~10 PFlop

~ ? PFlop

Software complexity

Model for UV processor

- Highly parallel – consider something achievable – NVIDIA promises 20 TFlop in 2 years – assume 50 TFlop
- Approximate analysis of ops/sample: 20000/calibration loop
- Processor: €1000, 5 calibration loops, 50% efficiency,
- each processor processes ~ 1 GB/s of data
- Requirement: 100 PFlop (AA) 2 EFlop (dishes)
- Buffer 1 hr of data therefore we need to buffer 7.2 TB in a fast store Assume €200 per TB.
- UV-processing cost: Cost = €2.5m (TB/s)

AA	< €10m
Dishes	~ €125m

The SKA Processing Challenge

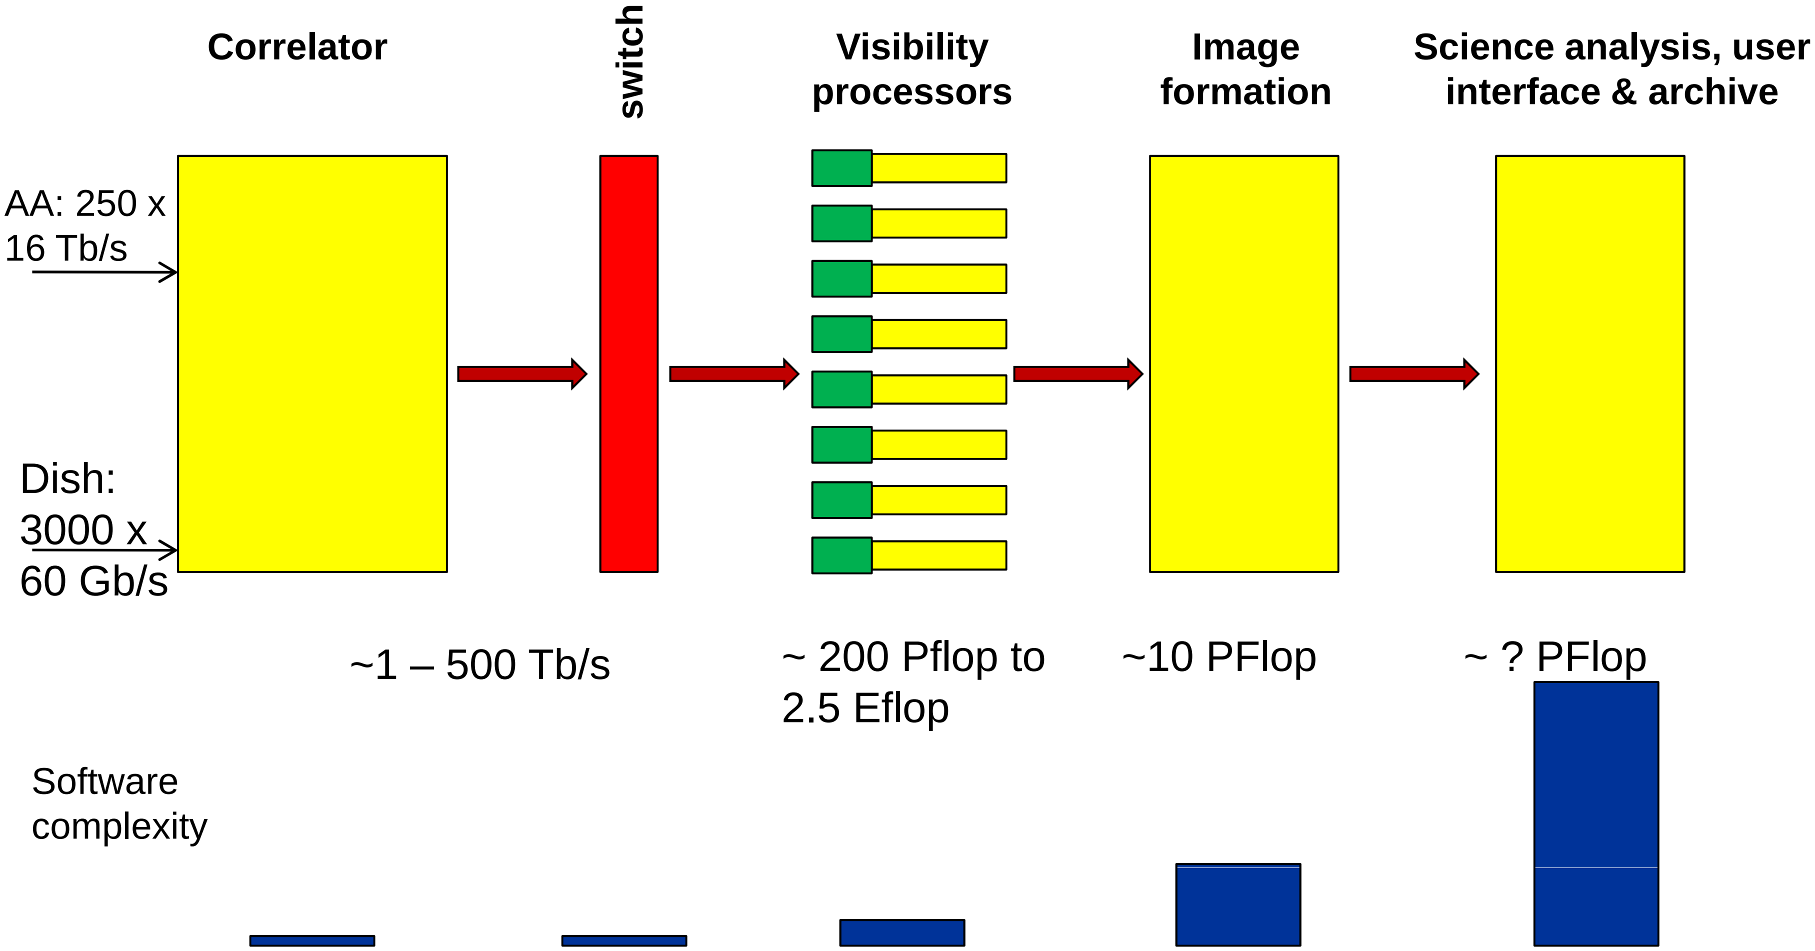

Correlator

switch

Visibility processors

Image formation

Science analysis, user interface & archive

AA: 250 x 16 Tb/s

Dish: 3000 x 60 Gb/s

~1 – 500 Tb/s

~ 200 Pflop to 2.5 Eflop

~10 PFlop

~ ? PFlop

Software complexity

Data Products

Experiment	T_{obs}	**B/km**	**D/m**	N_b	N_{ch}	N_v	**Size / TB**
High resolution spectral line	3600	200	15	1	32000	$5\ 10^{13}$	200
Survey spectral line medium resolution	3600	30	56	1000	32000	$8\ 10^{13}$	330
Snapshot continuum – some spectral information	60	180	56	1200	32	$7\ 10^{12}$	30
High resolution long baseline	3600	3000	60	1	4	$7\ 10^{14}$	360

- ~0.5 – 10 PB/day of image data
- Source count ~10^6 sources per square degree
- ~10^{10} sources in the accessible SKA sky, 10^4 numbers/record
- **~1 PB for the catalogued data**

100 Pbytes – 3 EBytes / year of fully processed data

Summary

- Processing challenge depends critically on overall system and experiment
- For dishes, the data rate drops significantly only after the UV processor
  - ☐ For many experiments the data rate can increase through the correlator.
- Design of processing system needs a complete analysis of the data flow
  - ☐ Separate UV-processor and imaging processors
  - ☐ Processing UV data is a highly parallel excellent scaleability
  - ☐ Local buffer for UV processor
- Maximum baseline and number of correlated elements drive costs
- Processing is expensive, but achievable
- **BUT must consider whole system – the processor MUST match the requirements of the system – PrepSKA**

Non PDRA resources

- Data flow and overview

  Alexander, Diamond, Garrington,

  ASTRON (de Vos, de Bruyn?)

- Pipeline processing:

  ASTRON LOFAR team

  CASU – Cambridge Astronomical Survey Unit – Irwin

  TDP – to be determined

- Data products and VO

  UK VO team – Walton, Richards; Astrowise Groningen

- Hardware architecture and implementation

  Centre for scientific computing (Cambridge),

  Oxford e-Science centre

  TDP - Kembal

Data from the correlator

Data rate depends on the experiment f = 1000 MHz.

Minimum data rate for dishes

- 0.8 square degrees, 2500 15-m dishes out to 200 km baseline Δf = 500MHz

$$\frac{R}{TSamples/s} = 6\left(\frac{\Delta f}{500MHz}\right)\left(\frac{B}{200km}\right)^{2}\left(\frac{D}{15m}\right)^{-2}\left(\frac{N_d}{2500}\right)^{2}\left(\frac{N_b}{1}\right)$$

High data rate for Aperture Array

- 75 square degrees, 250 56-m stations out to 200 km baseline Δf = 500MHz

$$\frac{R}{TSamples/s} = 6\left(\frac{\Delta f}{500MHz}\right)\left(\frac{B}{200km}\right)^{2}\left(\frac{D}{56m}\right)^{-2}\left(\frac{N_d}{250}\right)^{2}\left(\frac{N_b}{1500}\right)$$

>24 PB temporary buffer required

Processing

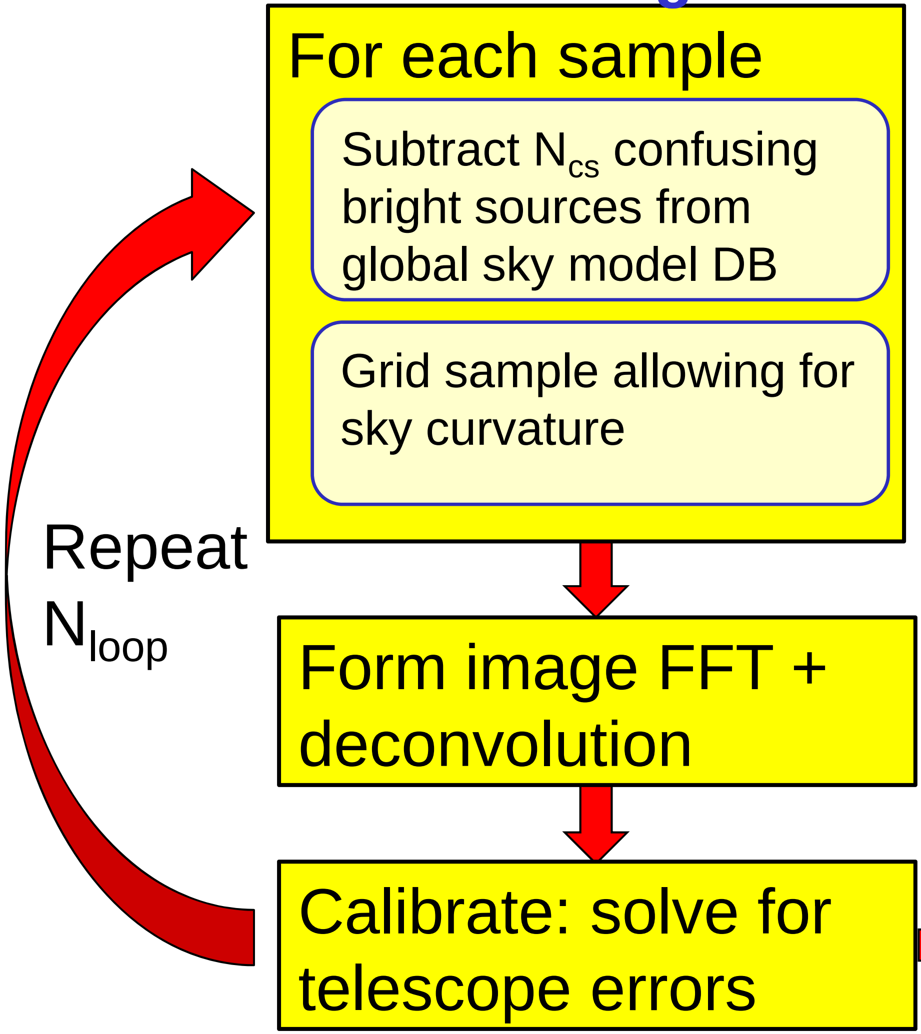

For each sample

Subtract N_{cs} confusing bright sources from global sky model DB

Grid sample allowing for sky curvature

Repeat N_{loop}

Form image FFT + deconvolution

Calibrate: solve for telescope errors

Update global sky model

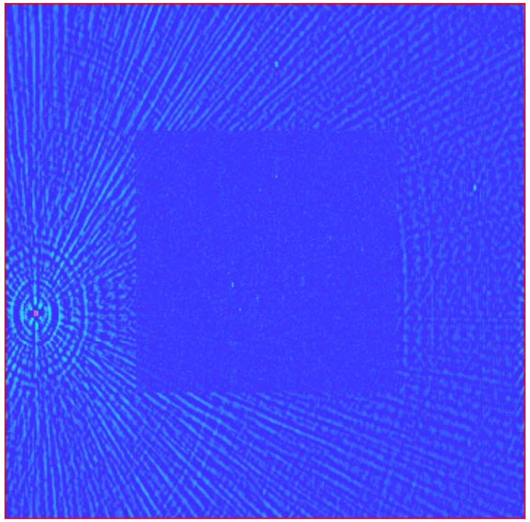

Effective ~20000 ops/sample using current algorithms

Store ~50TB/day of Observations

SKA Overall Structure