SKA Science Data Processor

B. Nikolic &
SKA SDP Consortium Team

Astrophysics Group, Cavendish Laboratory, University of Cambridge
http://www.mrao.cam.ac.uk/~bn204/

7 May 2013
NRAO Charlottesville

Aims

- SKA “Science Data Processor”
  - What it is?
  - The computing and data challenge
- Architectures and approaches to solving this challenge

- Similarities cf LLST? Differences?

Outline

Introduction

SKA_1-Mid

SKA_1-Low

SKA_1-Survey

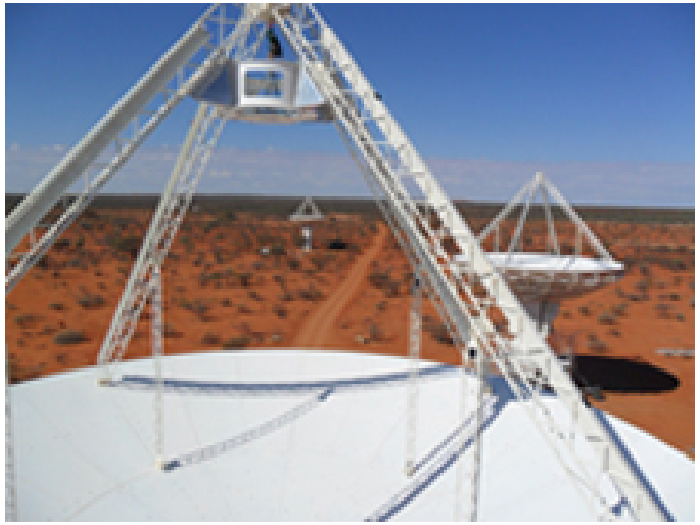

Introduction

SKA_1-Mid

SKA_1-Low

SKA_1-Survey

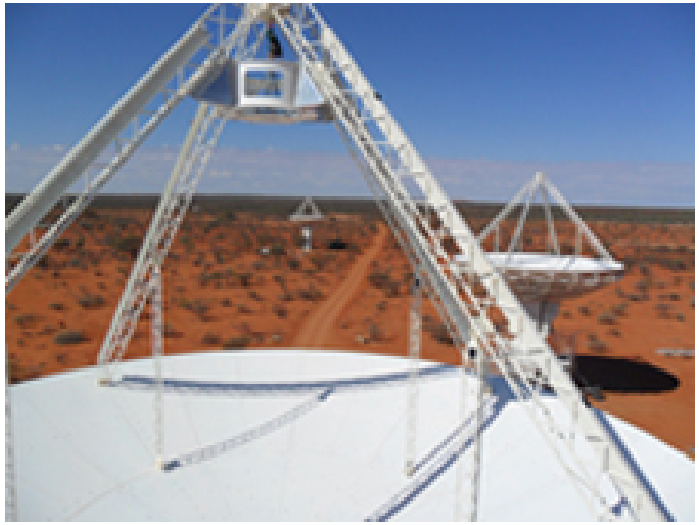

HPC Computer

Introduction

SKA_1-Mid

SKA_1-Low

SKA_1-Survey

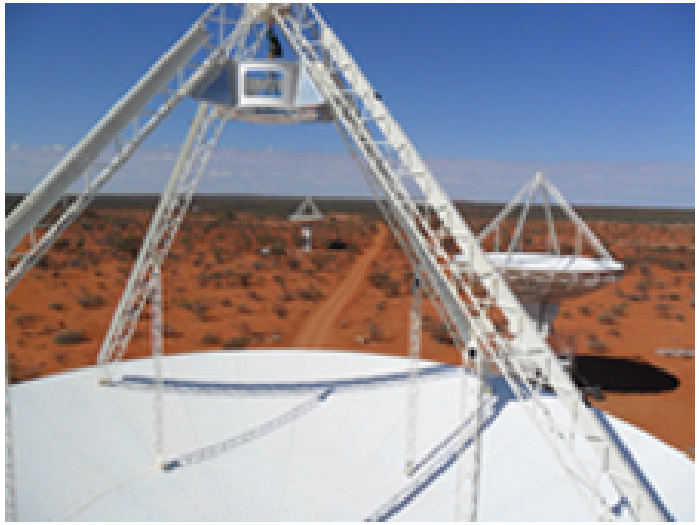

HPC Computers

SKA

Summary of planned features:

- 100$\times$ better sensitivity than current best telescopes
- $10^6\times$ sky survey speed compared to current facilities
- Frequency coverage 50 MHz–$\sim$ 10 GHz
- Extremely radio quite sites
- Staged construction:
  - Precursors: ASKAP/MeerKAT – under construction
  - SKA_1: Construction start in 2016-7. Full Ops 2020
  - SKA_2: Construction start in $\sim$2023
- Will require 'exascale' computing to form images and analyse them

SKA_1 Sites & Receptors

Karoo Desert, South Africa

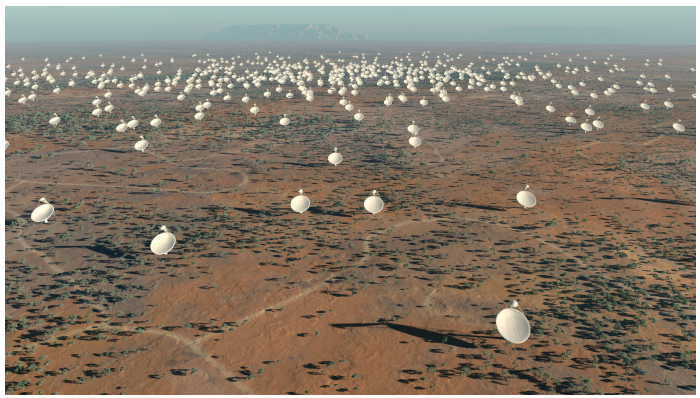

Western Australia

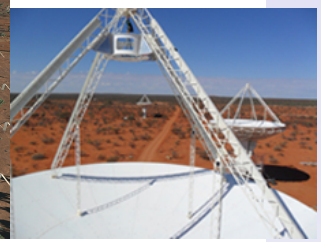

SKA as an ICT project

Within the SKA project, Information and Communications Technology (ICT) is:

- Major part of the design budget
- Major part of the construction budget
- Major part of the operations budget (power, S/W maintenance)
- Major part of the risk 'budget'
  - Strong interaction with industry, other astronomy projects, other science project
  - Direct involvement of major existing HPC facilities
  - Reusing, Road-mapping, Best-practices, Sustainability

Outline

Wide area dataflow for SKA

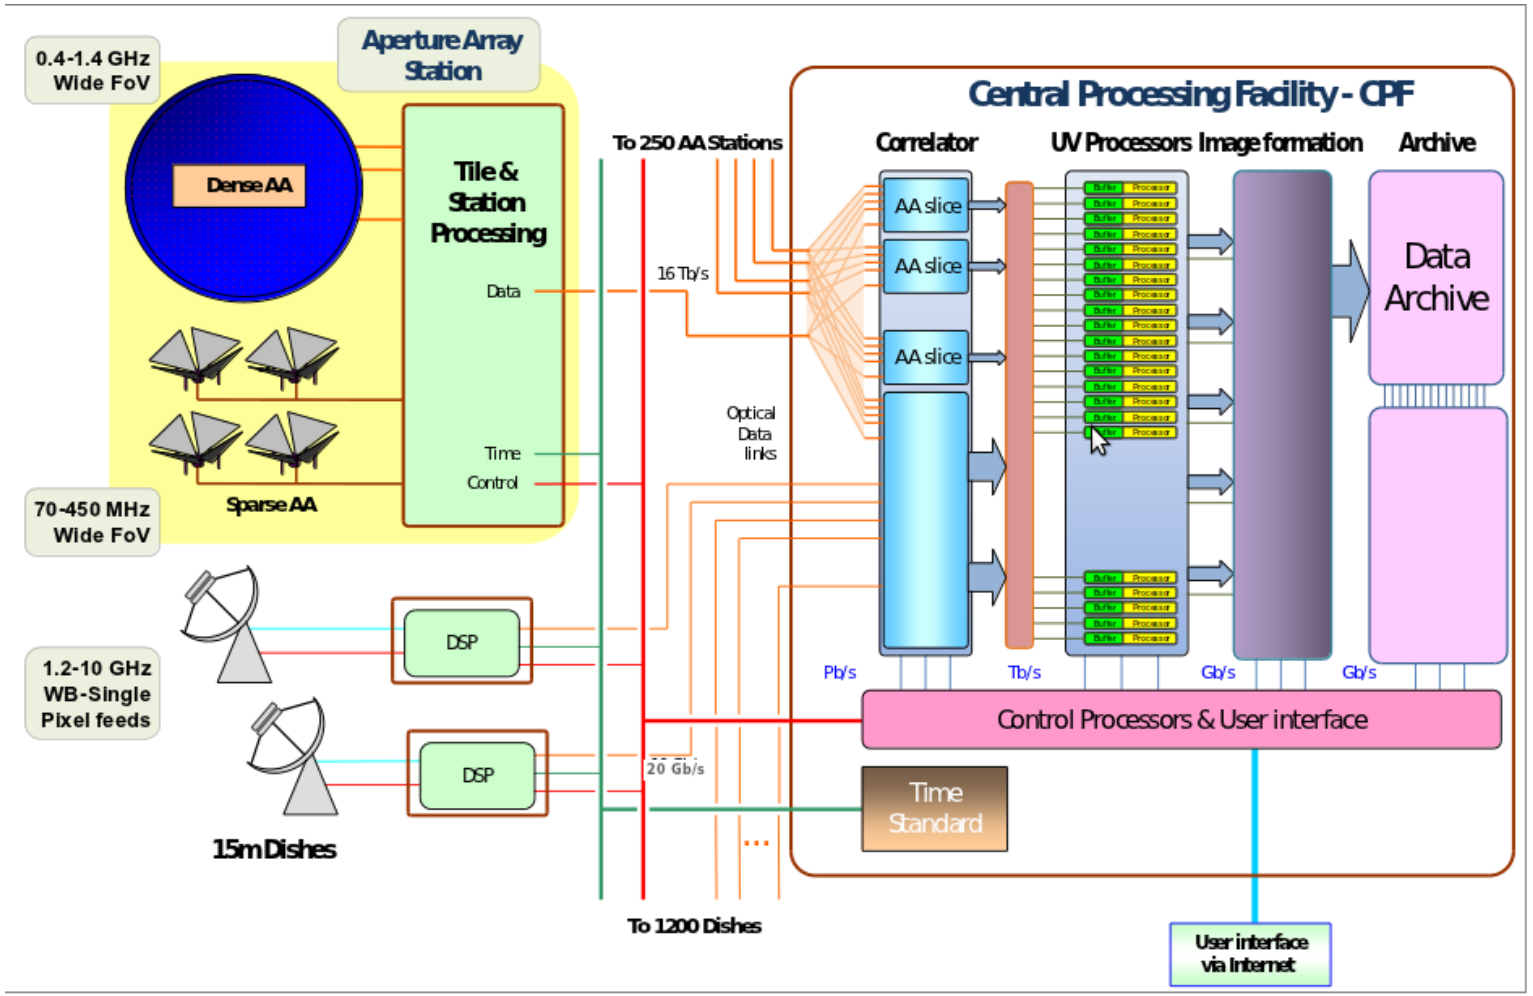

SKA simplified dataflow

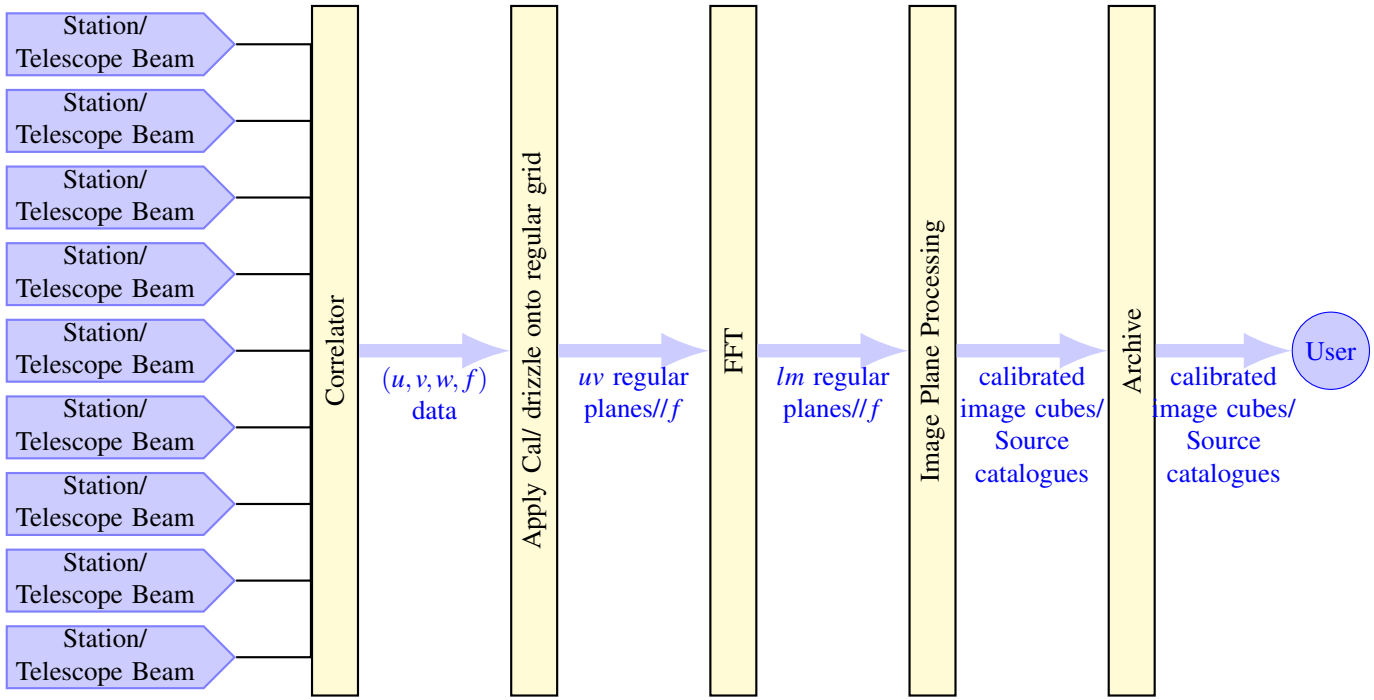

What drives the challenge?

Complex to do

- Irregular 3d sampling of the signal
  $\Rightarrow$ grid & expensive 3d correction
- Image reconstruction (CLEANing)
  $\Rightarrow$ identify and remove 'bright' sources
- Changing electrical properties of telescope and the atmosphere
  $\Rightarrow$ self-calibrate

Computationally intensive

- Very high data rates

Calibration/CLEAN

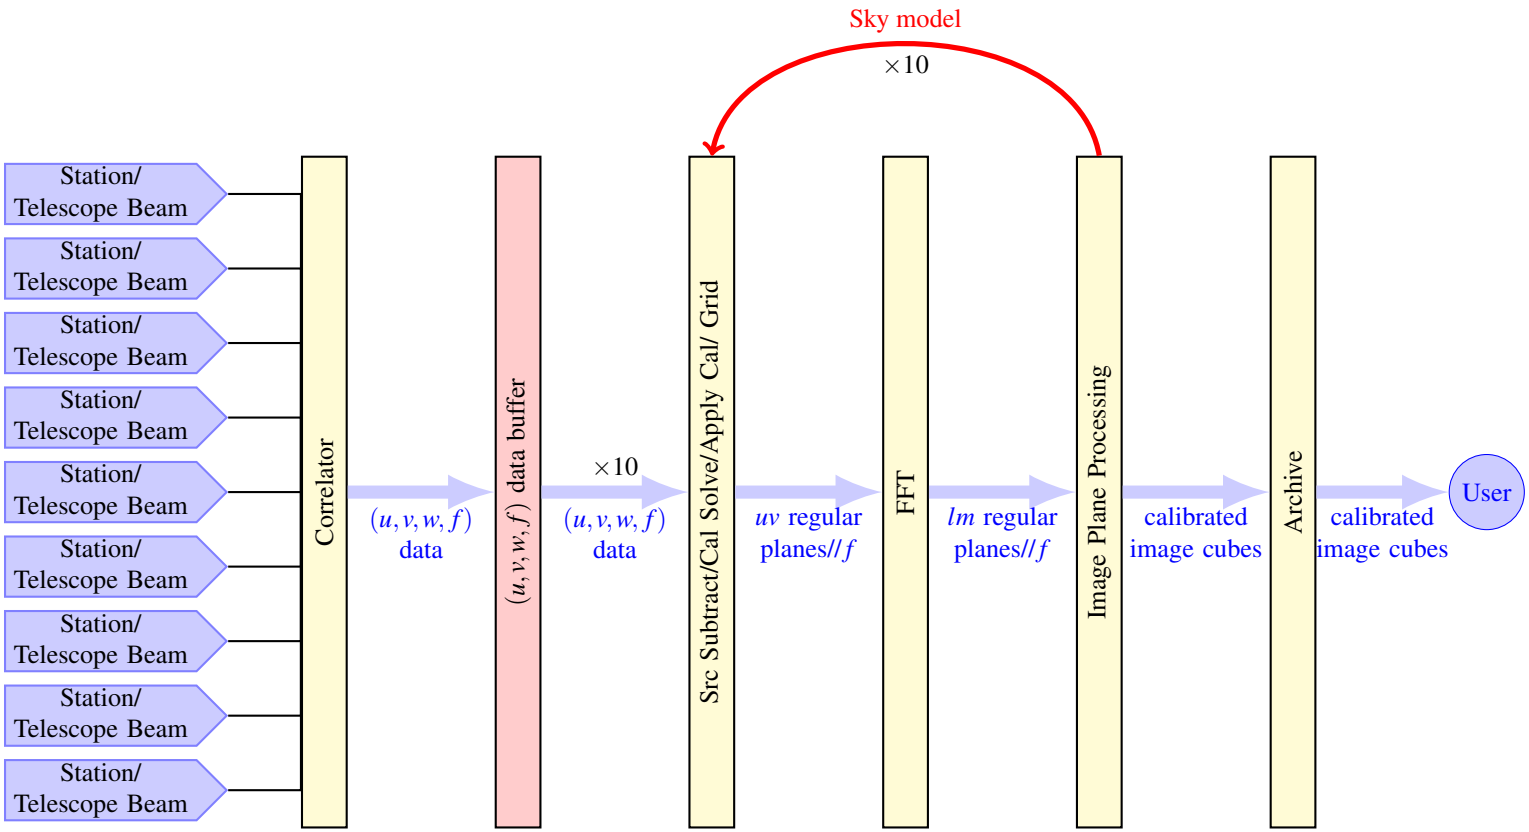

Data rate – high N

- # visibilities scales as N^2
- Very large N:
  - $N \sim 900$ for SKA_1-Low
  - $N \sim 250$ for SKA_2-Mid

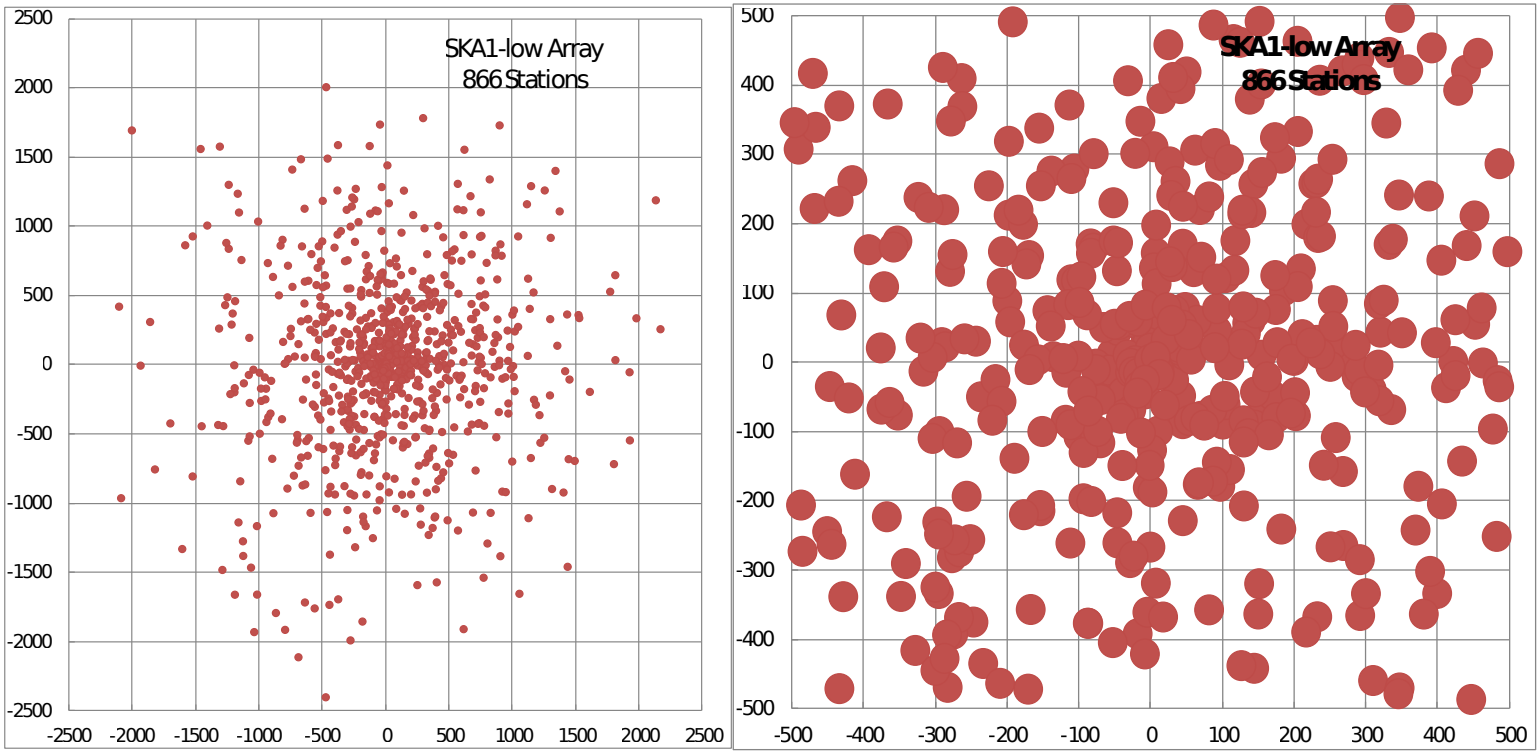

SKA_1-Low Core stations configuration. Credit: P. Dewney, SKA1 Baseline Design/RFP

Data rate – multiple beams

- # visibilities scales as # beams
- SKA_1-Survey will have 36 simultaneous beams

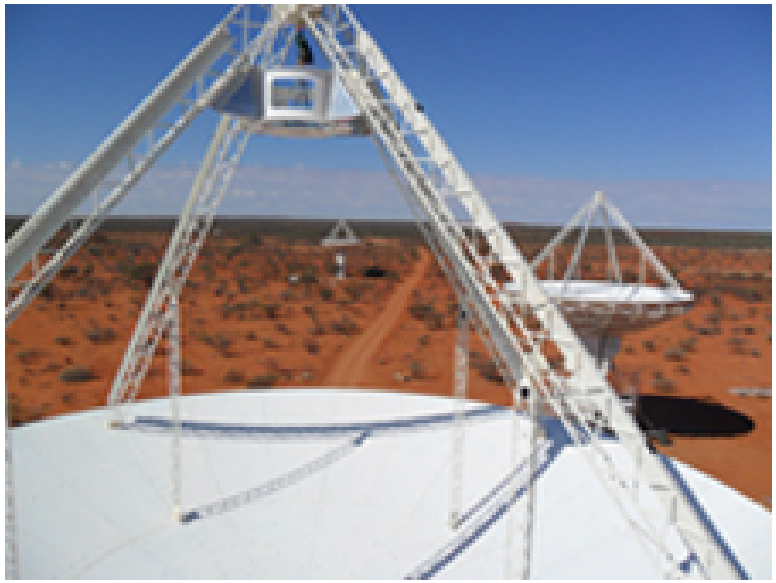

Data Rate – time smearing

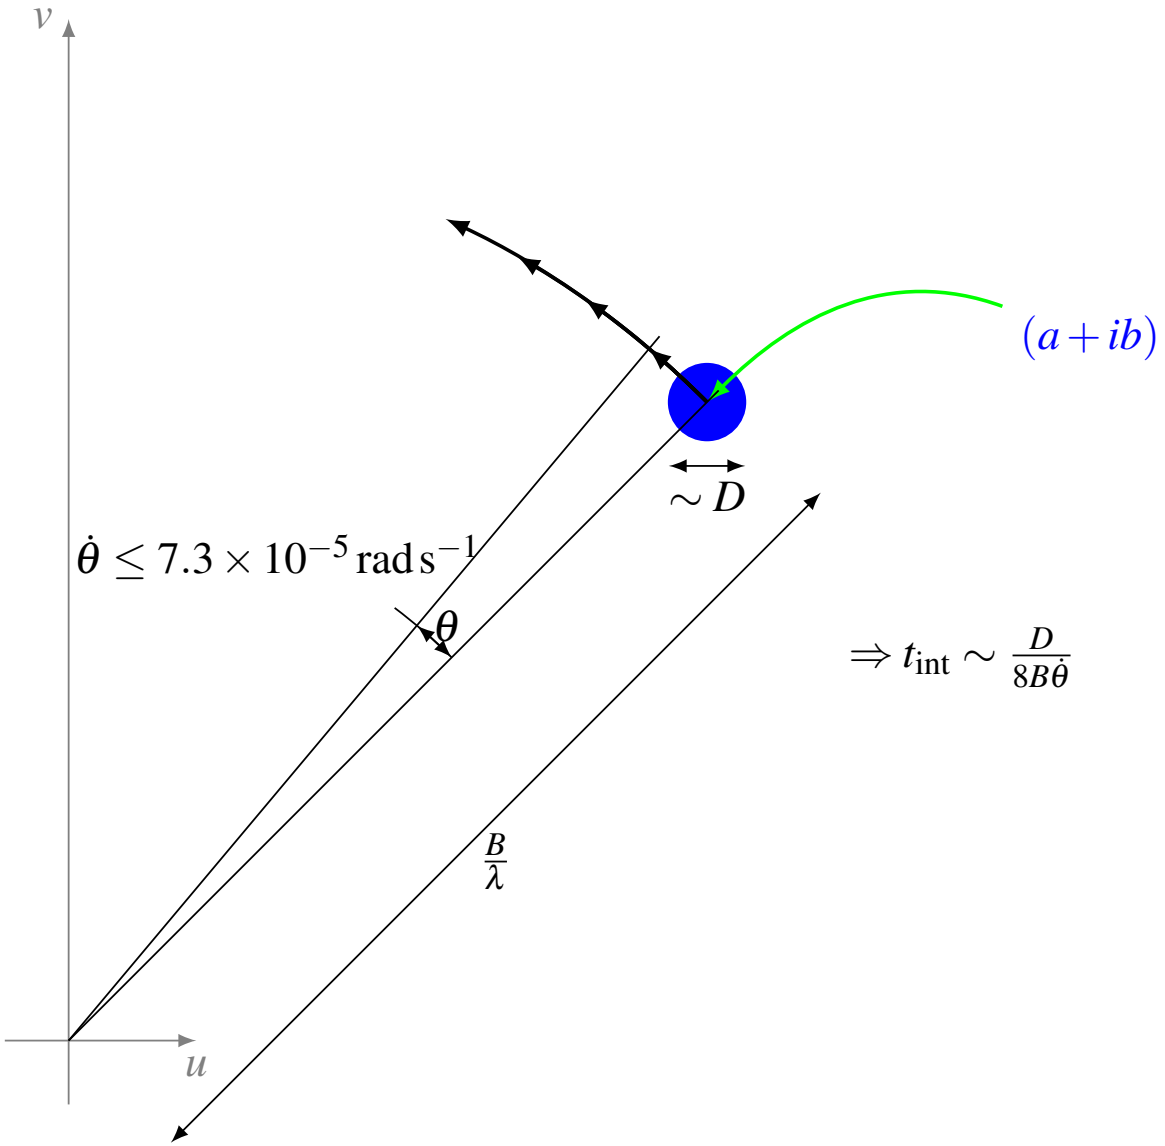

Data Rate – bandwidth smearing

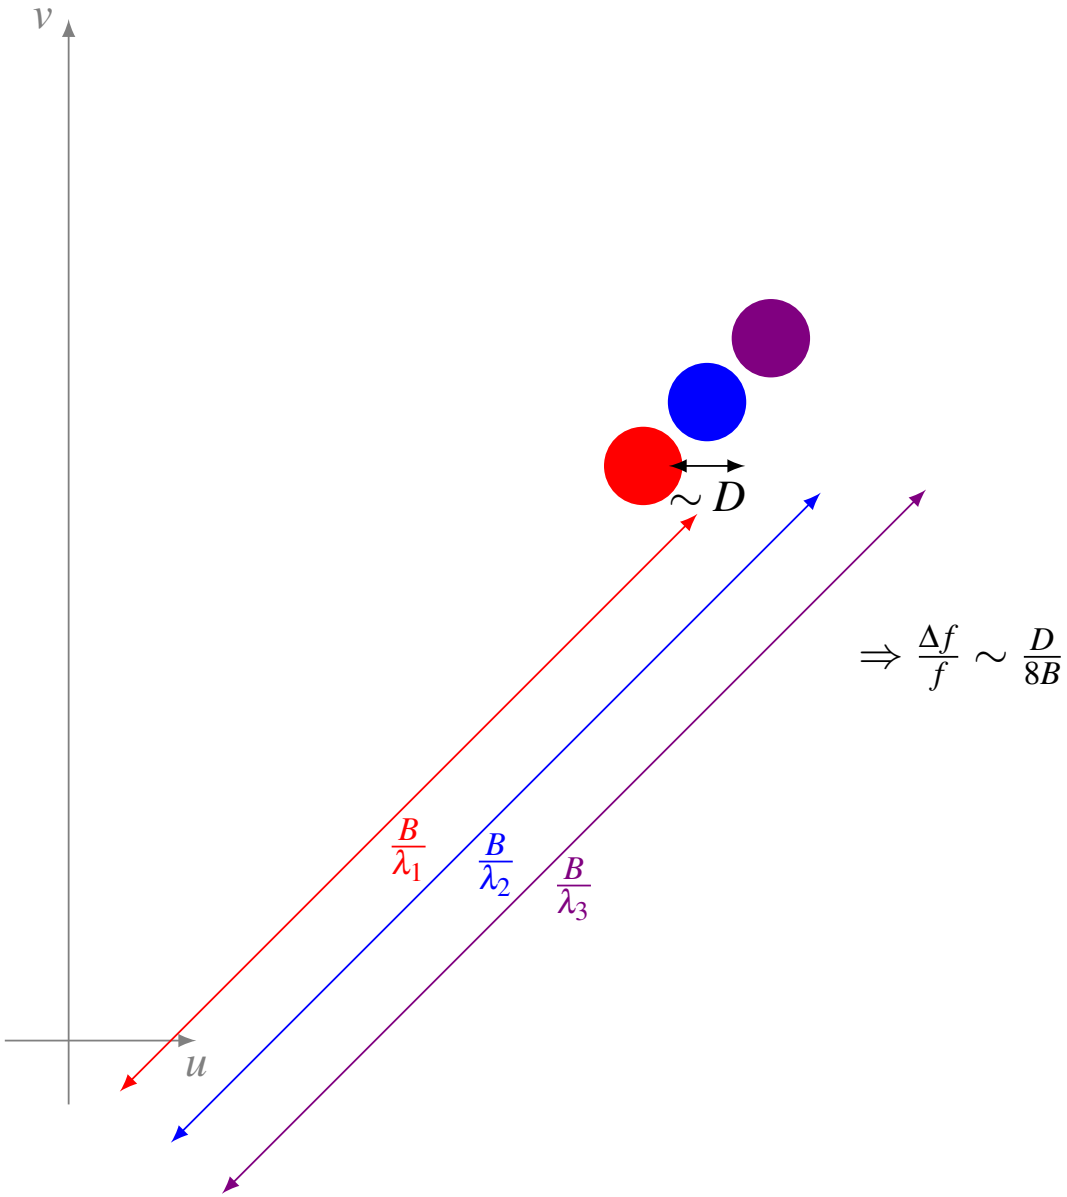

$$\Rightarrow \frac{\Delta f}{f} \sim \frac{D}{8B}$$

Data Products

- Data volume most demanding for 'high resolution spectral line survey' observation:
  - Roughly 32k$\times$32k$\times$32k cube
    $\Longrightarrow \sim 200\,\mathrm{TB}$/obs
  - Very low re-observation cadence
  - Potentially useful for cosmology $\rightarrow$ many fields?
- Other observations typically coarser spatial and/or frequency resolution $\rightarrow$ TB scale datasets
- However some of the other observations could have high re-observation cadence
- Transient Imaging pipeline for selectively retaining data

Data flow with rates, computes and storage requirements

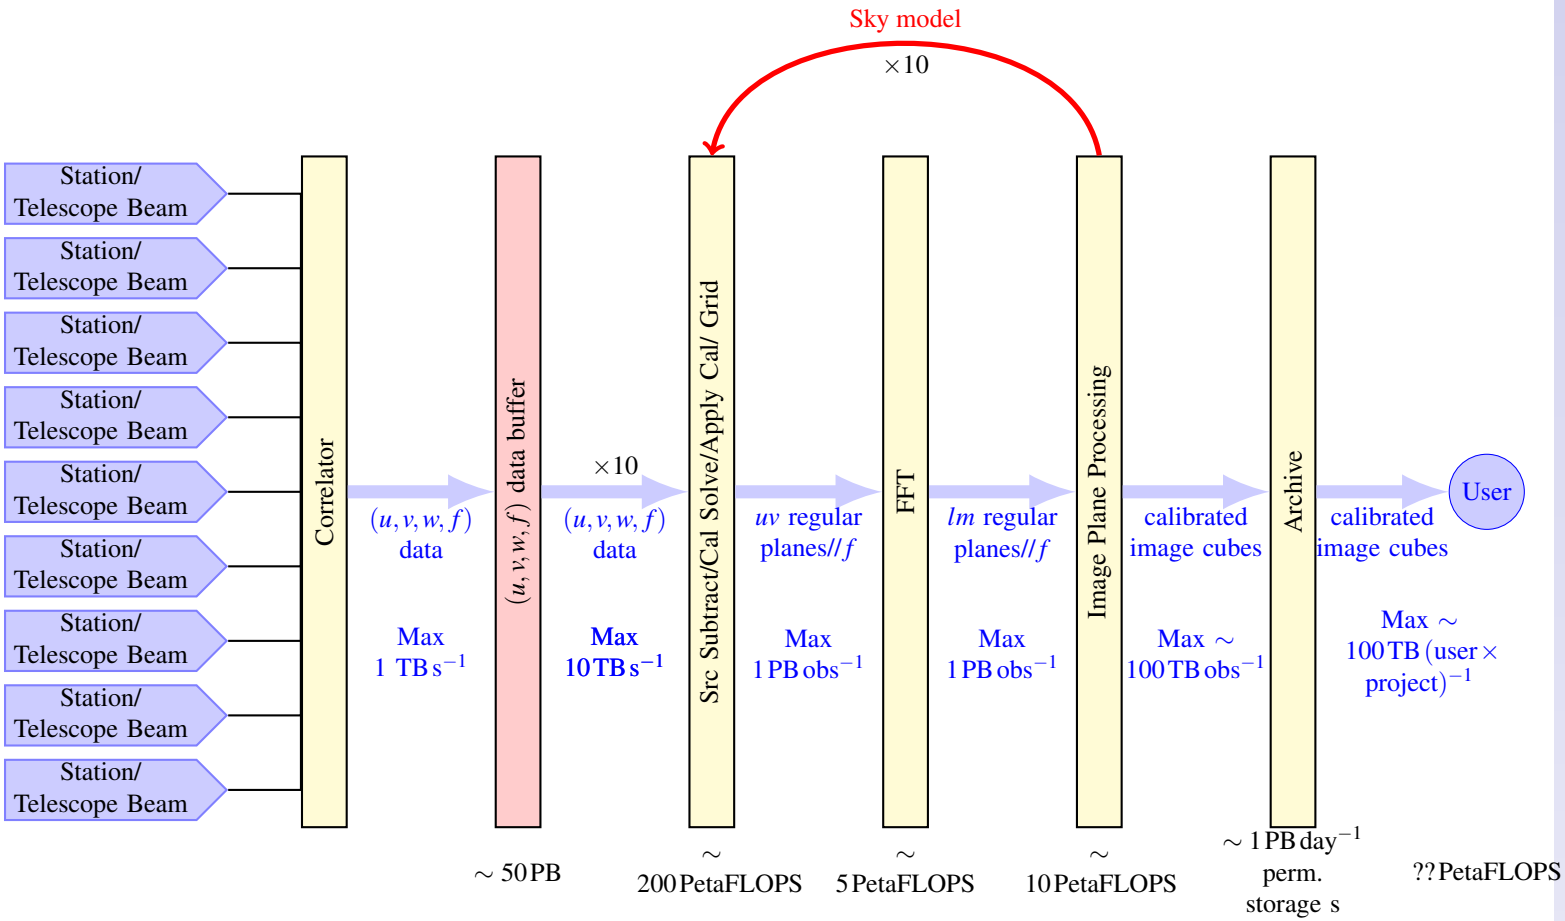

Outline

Technology

Rapid progress, often driven by consumer technologies:

- $> 10\,\mathrm{TeraFLOPS}/$ accelerator card likely in 2017
  - $100\,\mathrm{PetaFLOPS} \implies 10^4 \mathrm{cards}, \sim 2\,\mathrm{MW}$
  - At 2 cards/per node, $> 50\,\mathrm{TB\,s^{-1}}$ total inter-node network capacity easily achieved
- $> 1\,\mathrm{TB\,s^{-1}}$ total network throughput possible today (600 port Infiniband FDR)
- Nodes with $2\times$FDR $+2 \times 40\,\mathrm{GbE}$ already being deployed
  $\implies 10/\mathrm{GB\,s^{-1}}$ onto one node
- $40\,\mathrm{kW}/\mathrm{rack}$ with water-cooled cabinet doors
  - $100\,\mathrm{PetaFLOPS}, 2\,\mathrm{MW} \implies 50\,\mathrm{racks}$

High-level dataflow/possible architecture

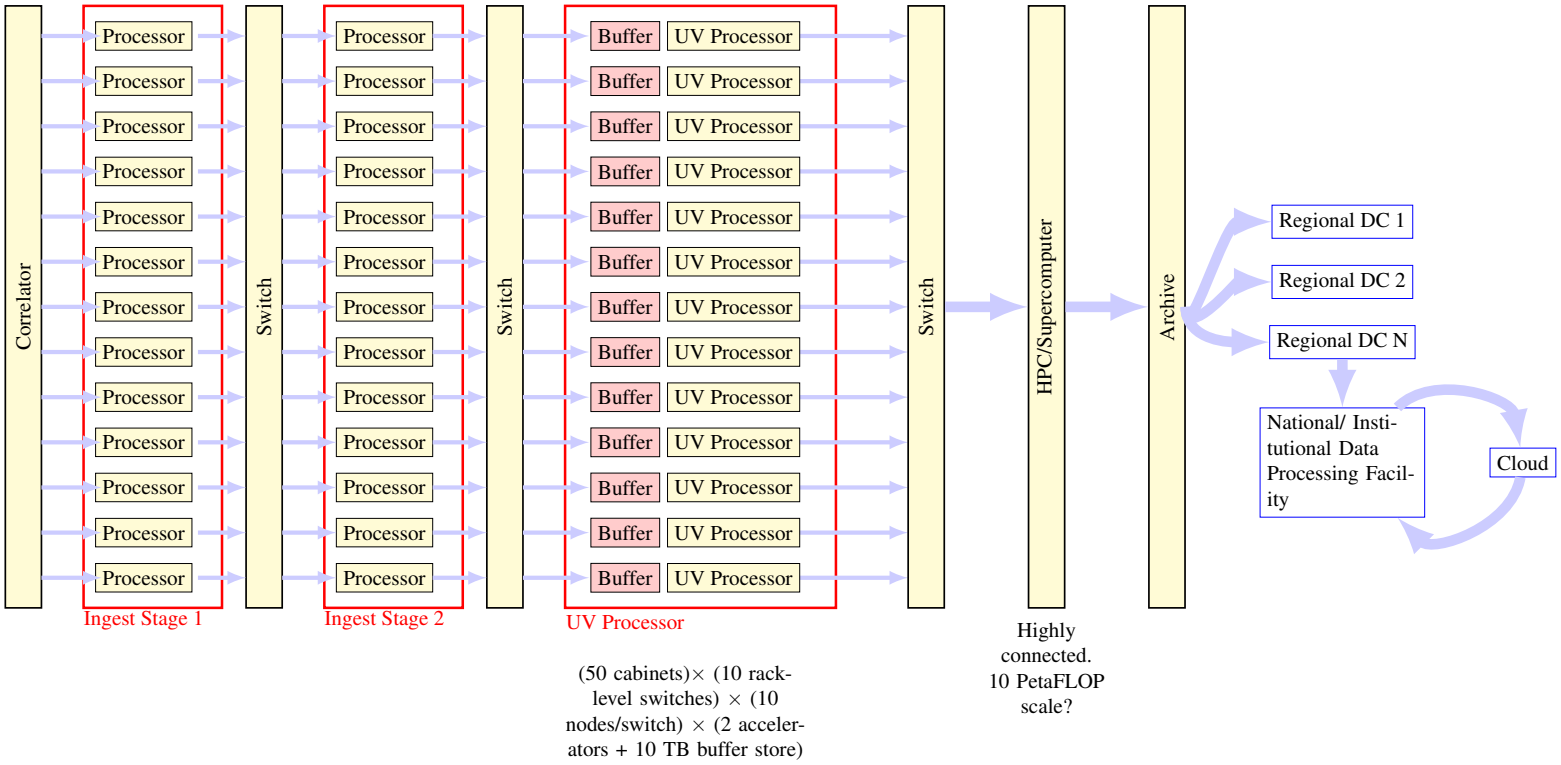

Data Distribution, Remote Compute & Visualisation

- Tiered data distribution allows use of existing/forthcoming national compute facilities
- Remote (institutional/national/international) visualisation and compute
- Efficient collaboration for large science teams essential

CyberSKA

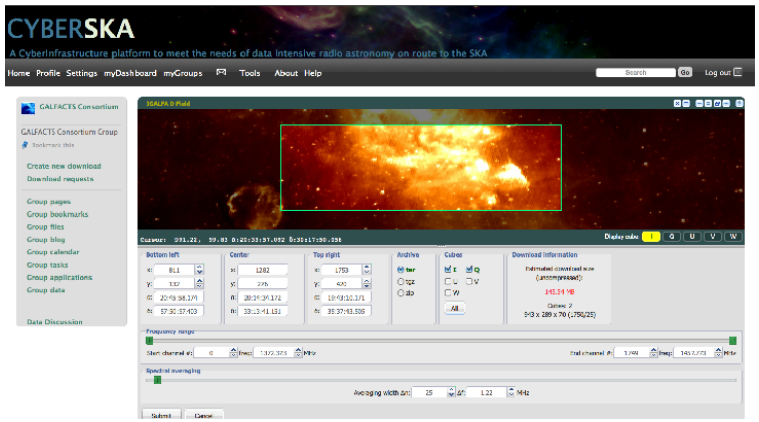

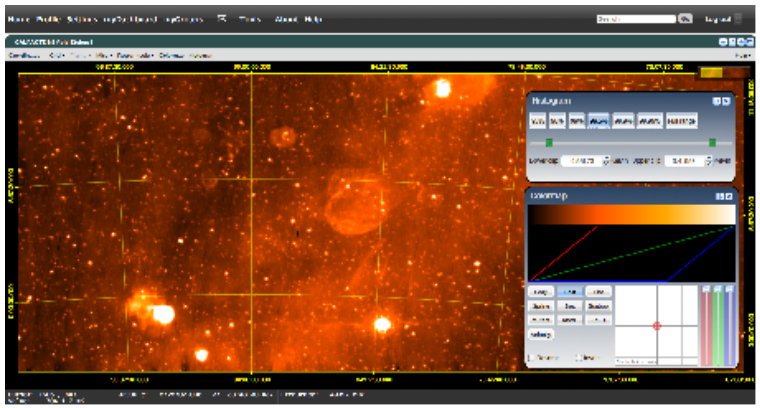

Software

- A large software engineering project: hybrid between conventional HPC, ‘big data’ and streaming processing
- Need both flexibility and the ability to operate autonomously (without human intervention)
- Data flow and data-locality critical
- Few *iterations*, mostly data-parallel
- Streaming / soft-realtime processing
- Evolving underlying hardware

Outline

Summary

- Very large data rates in SDP:
  - Internal rates $\sim 1 - 10\,\mathrm{TB\,s^{-1}}$
  - Into archive $\sim 0.1 - 1\,\mathrm{PB\,day^{-1}}$
- Complex, varied, processing *before* data reduction
- SDP cost analysed and balanced against overall SKA system
- Commercial COTS H/W roadmaps look promissing
- Challenges:
  - S/W design, construction, maintenance
  - System complexity
  - H/W & S/W failures
  - Schedule
  - Science Analysis on large data-cubes