The Science Data Processor and Regional Centre Overview

Paul Alexander

UK Science Director the SKA Organisation
Leader the Science Data Processor Consortium

Welcome to this UK-SKA Meeting

For those of you who have not been to Cambridge for a few years you will notice some big changes, you are now in our new Astronomy Campus housing:

- IoA
- Cavendish Astrophysics
- Kavli Institute for Cosmology in Cambridge

New Battcock Centre for Experimental Astrophysics

SKA: A Leading Big Data Challenge for 2020 decade

SQUARE KILOMETRE ARRAY
SDP
SCIENCE DATA PROCESSOR

Antennas

Digital Signal Processing (DSP)

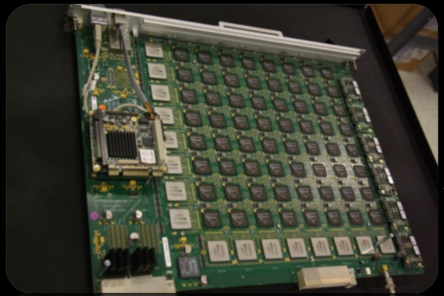

Transfer antennas to DSP
2020: 5,000 PBytes/day
2030: 100,000 PBytes/day

Over 10's to 1000's kms

To Process in HPC
2020: 50 PBytes/day
2030: 10,000 PBytes/day

Over 10's to 1000's kms

High Performance Computing Facility (HPC)

HPC Processing
2020: 300 PFlop
2028: 30 EFlop

UNIVERSITY OF CAMBRIDGE

SDP Organisation

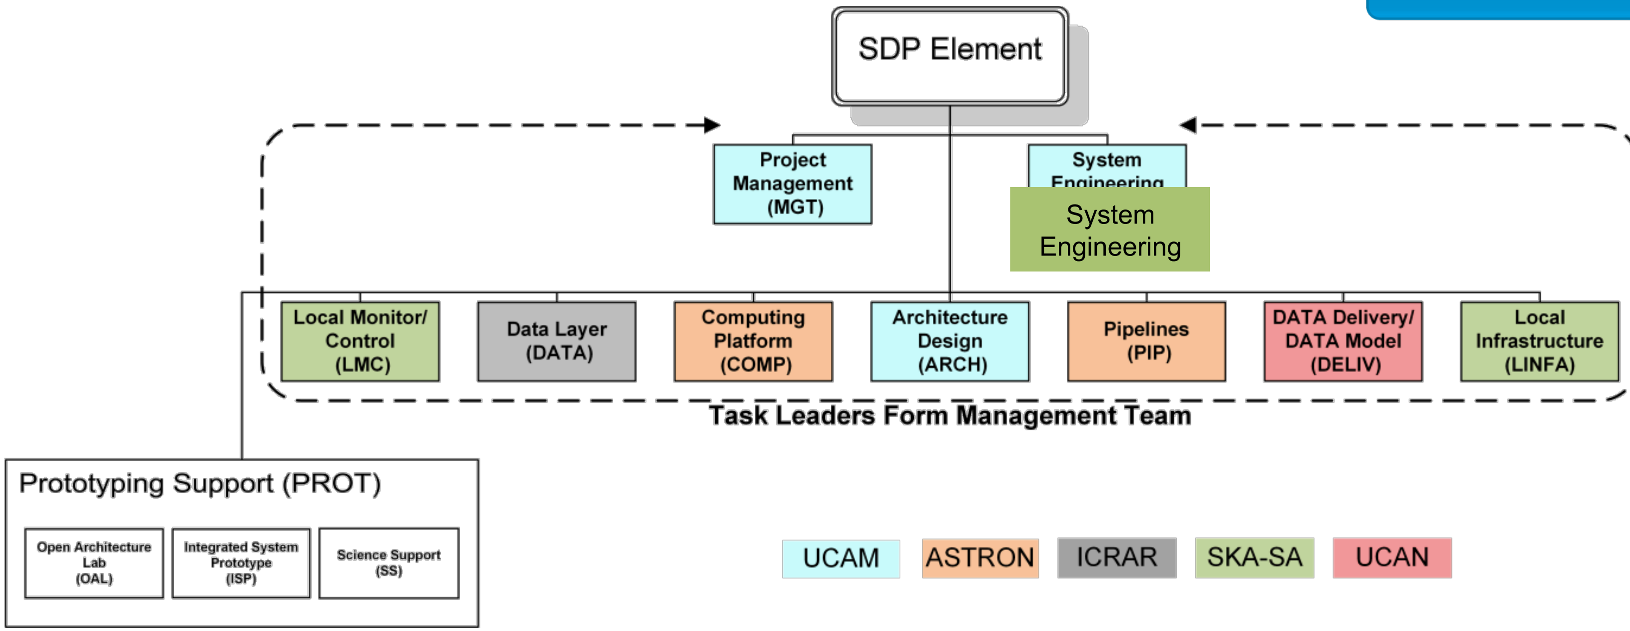

- **Lead:** **Paul Alexander**
- **PM:** **Jeremy Coles**
- Deputy PM: Ian Cooper
- **PE/Architect:** **Bojan Nikolic**
- SE: Ferdl Graser
- **PS:** **Rosie Bolton**

- COMP: Chris Broekema
- PIP: Ronald Nijboer
- DATA: Andreas Wicenec
- DELIV: Rob Simmonds
- LMC: Simon Ratcliffe
- LINF: Jasper Horrel

- **PIP-Imaging:** **Anna Scaife**
- **PIP-NIP:** **Ben Stappers**

UNIVERSITY OF CAMBRIDGE

Standard interferometer

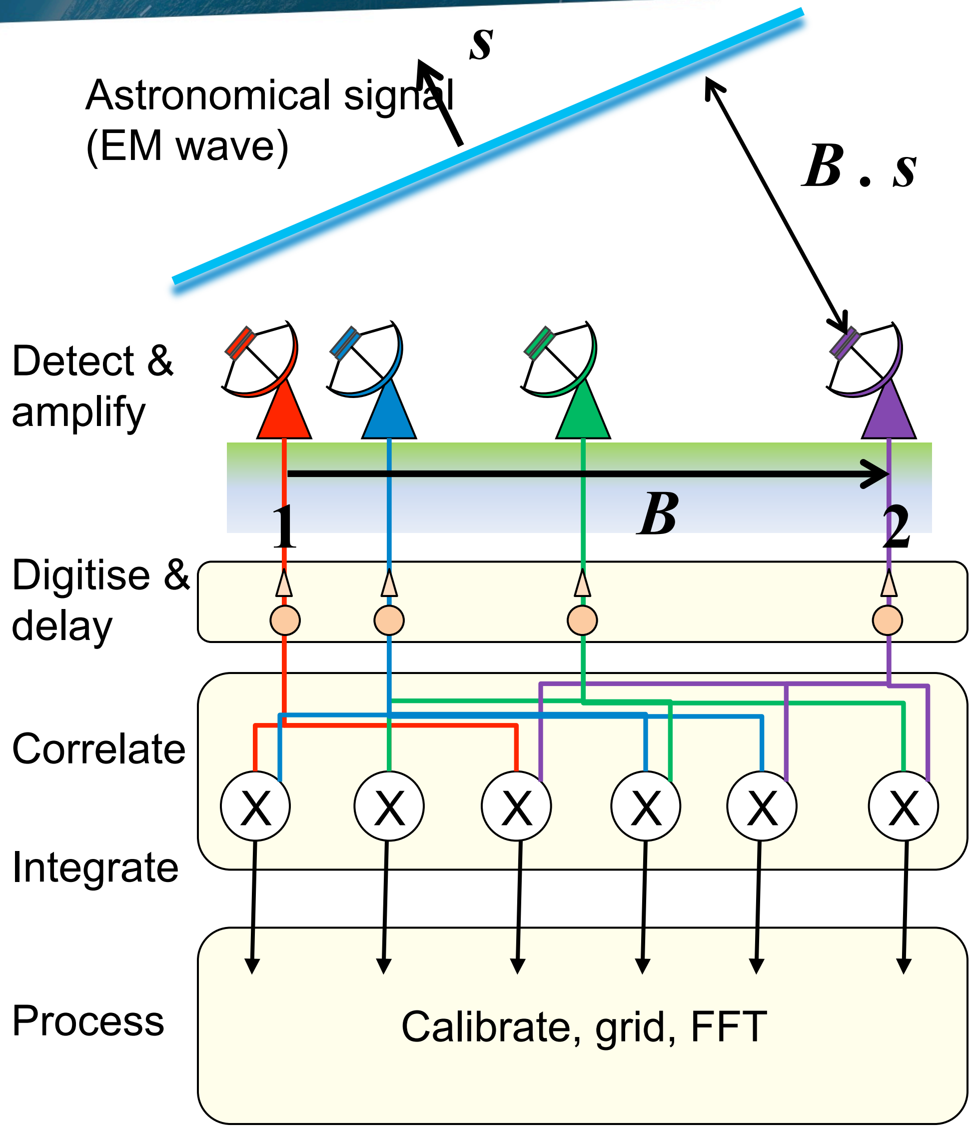

- Visibility:

$$V(\boldsymbol{B}) = E_1 E_2^* = I(\boldsymbol{s}) \exp(\, i\, \omega\, \boldsymbol{B}.\boldsymbol{s}/c\,)$$

- Resolution determined by maximum baseline
  $\theta_{max} \sim \lambda / B_{max}$

- Field of View (FoV) determined by the size of each dish
  $\theta_{dish} \sim \lambda / D$

SKY Image

Challenge Very Dependent on Experiment

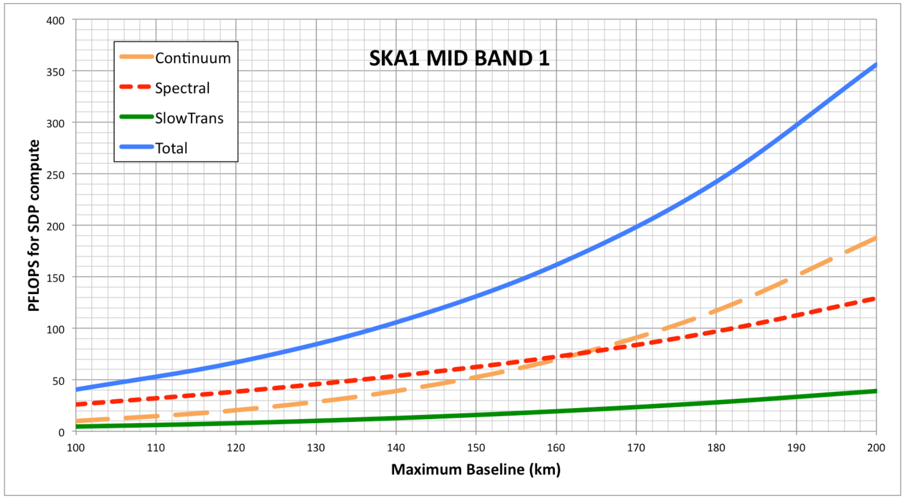

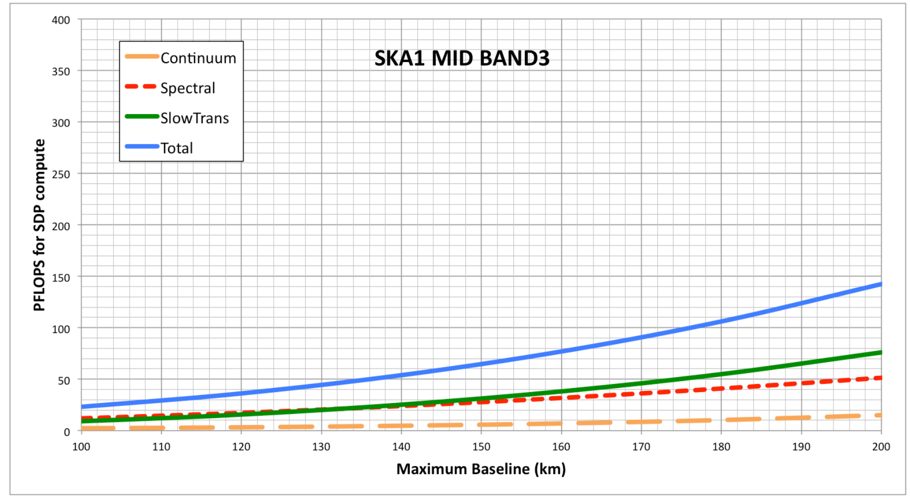

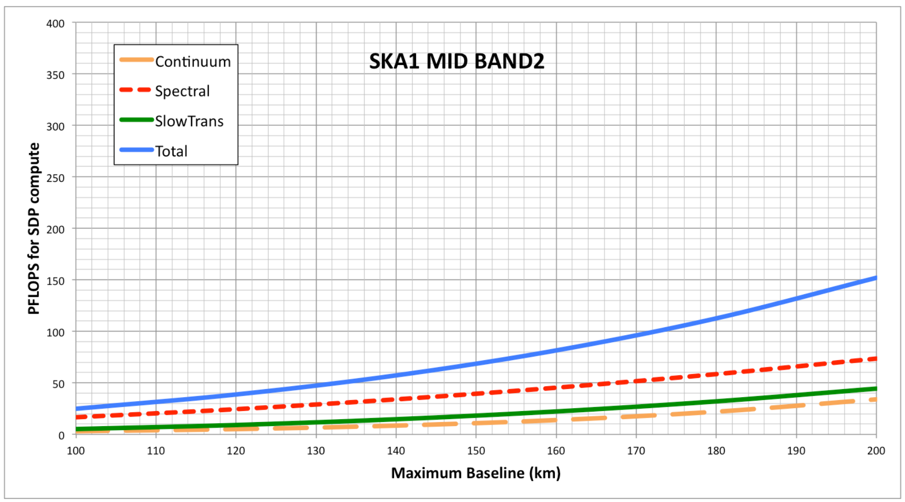

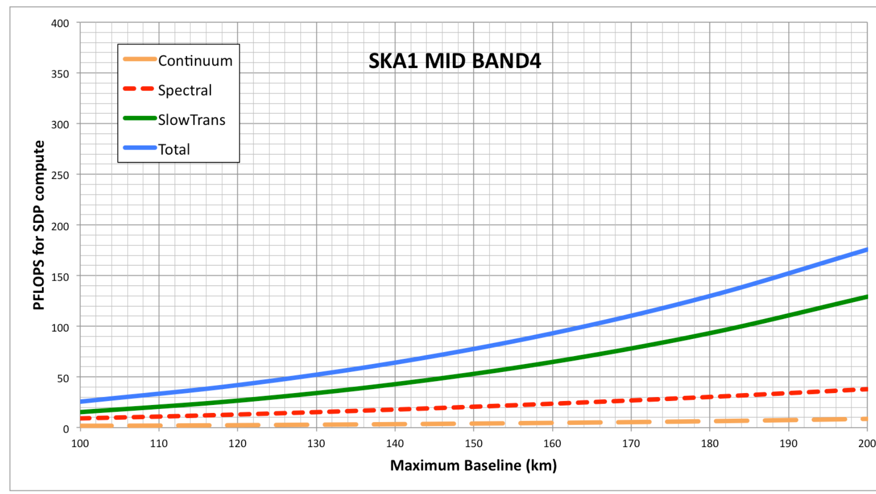

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

UNIVERSITY OF CAMBRIDGE

High Level Description

	Ingest (GB/s)
SKA1_Low	500
SKA1_Mid	1000

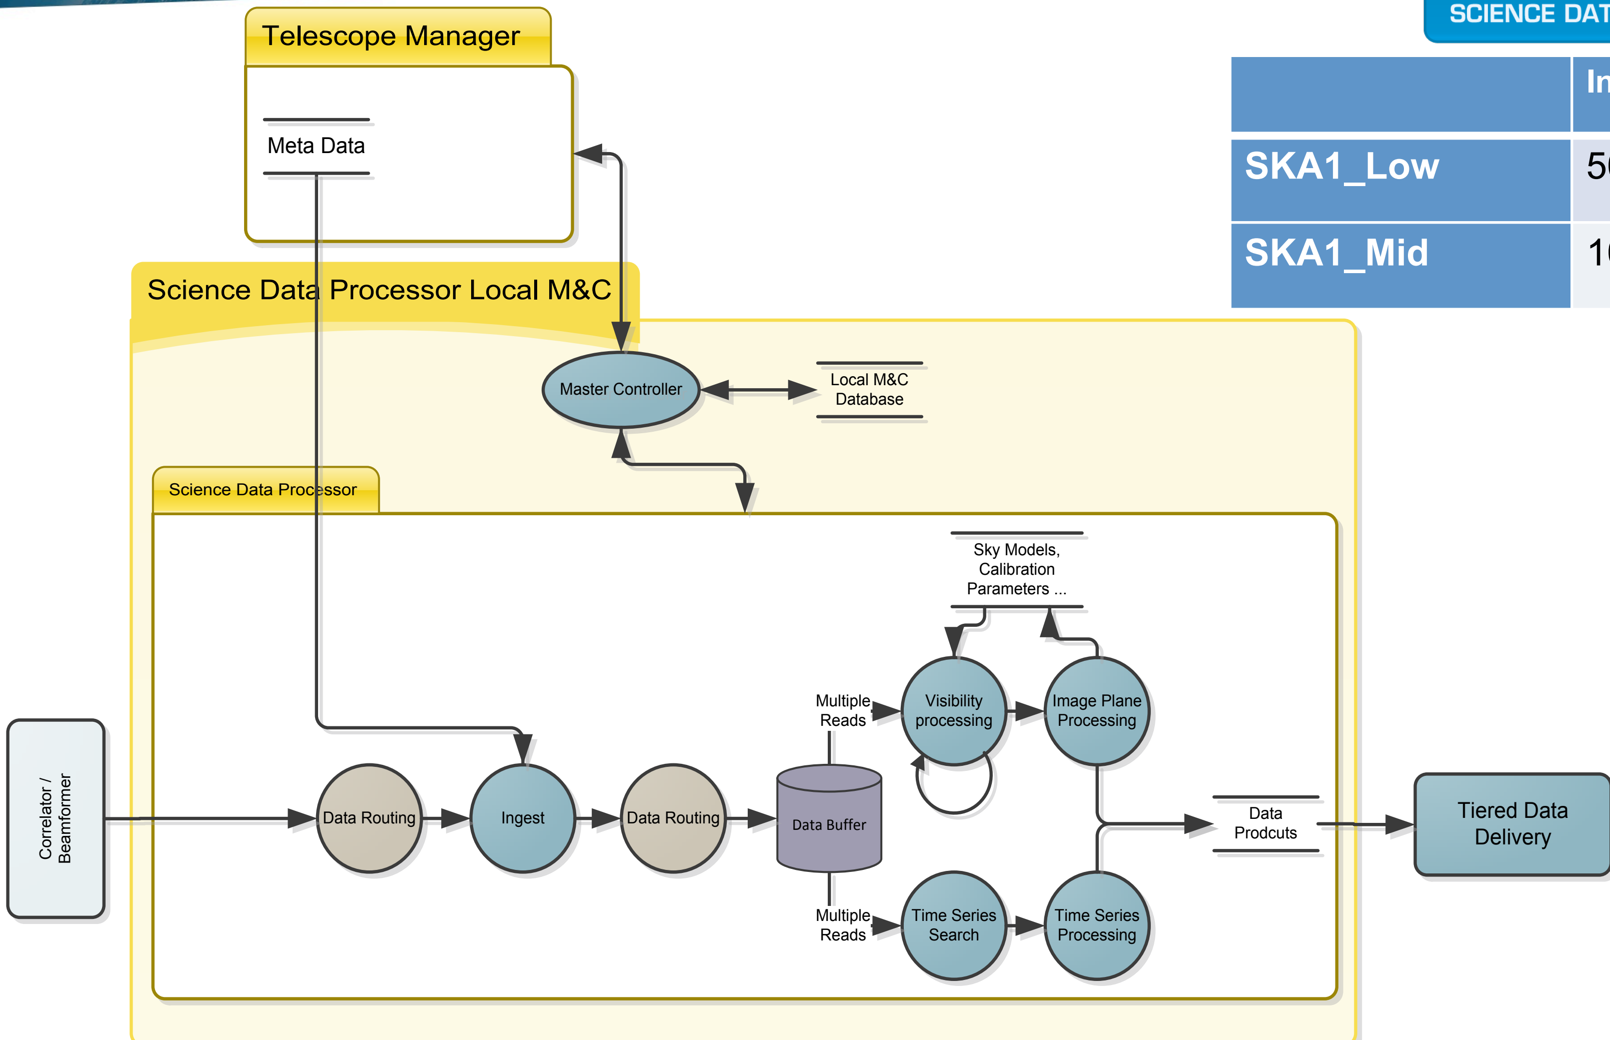

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

UNIVERSITY OF CAMBRIDGE

One SDP Two Telescopes

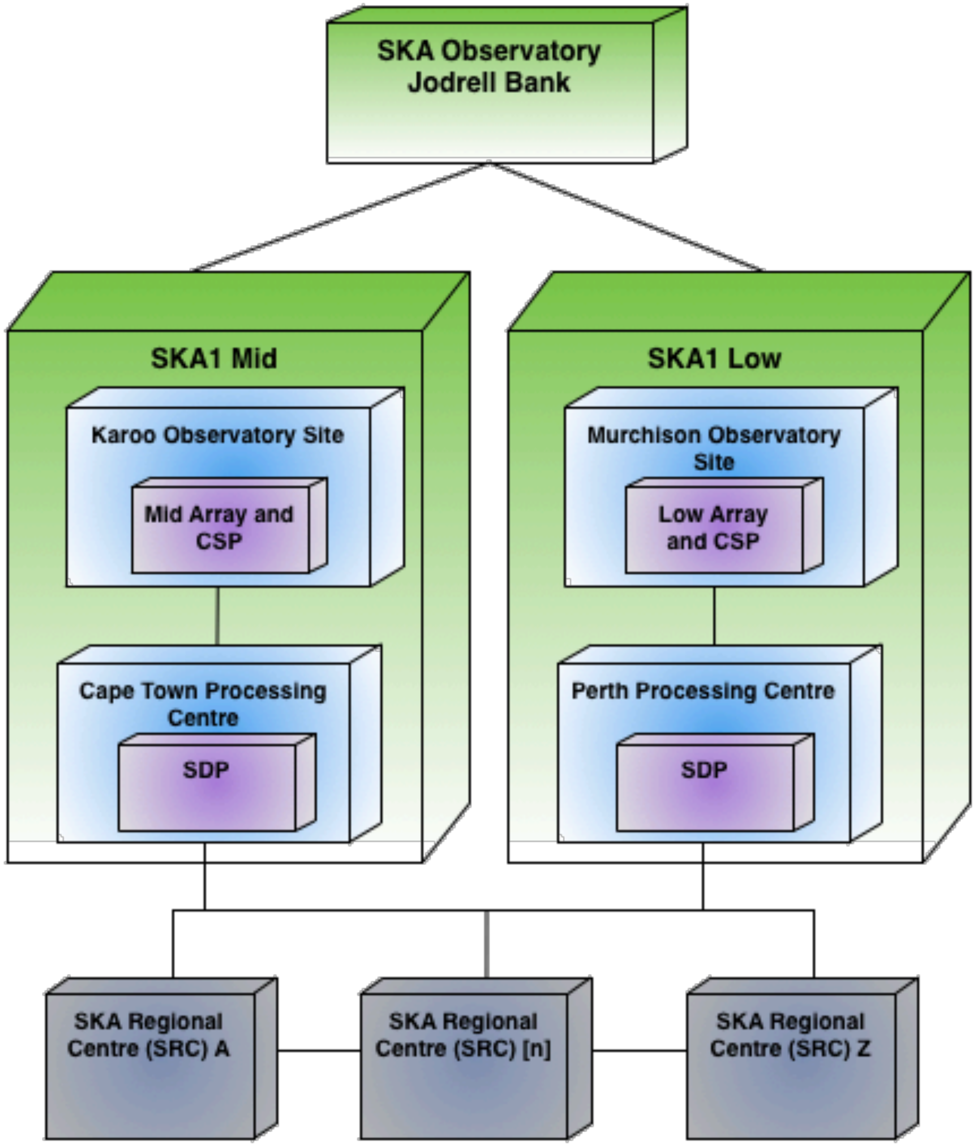

	Ingest (GB/s)
SKA1_Low	500
SKA1_Mid	1000

In total need to deploy eventually a system which is close to 0.5 EFlop of processing

… and regional centres

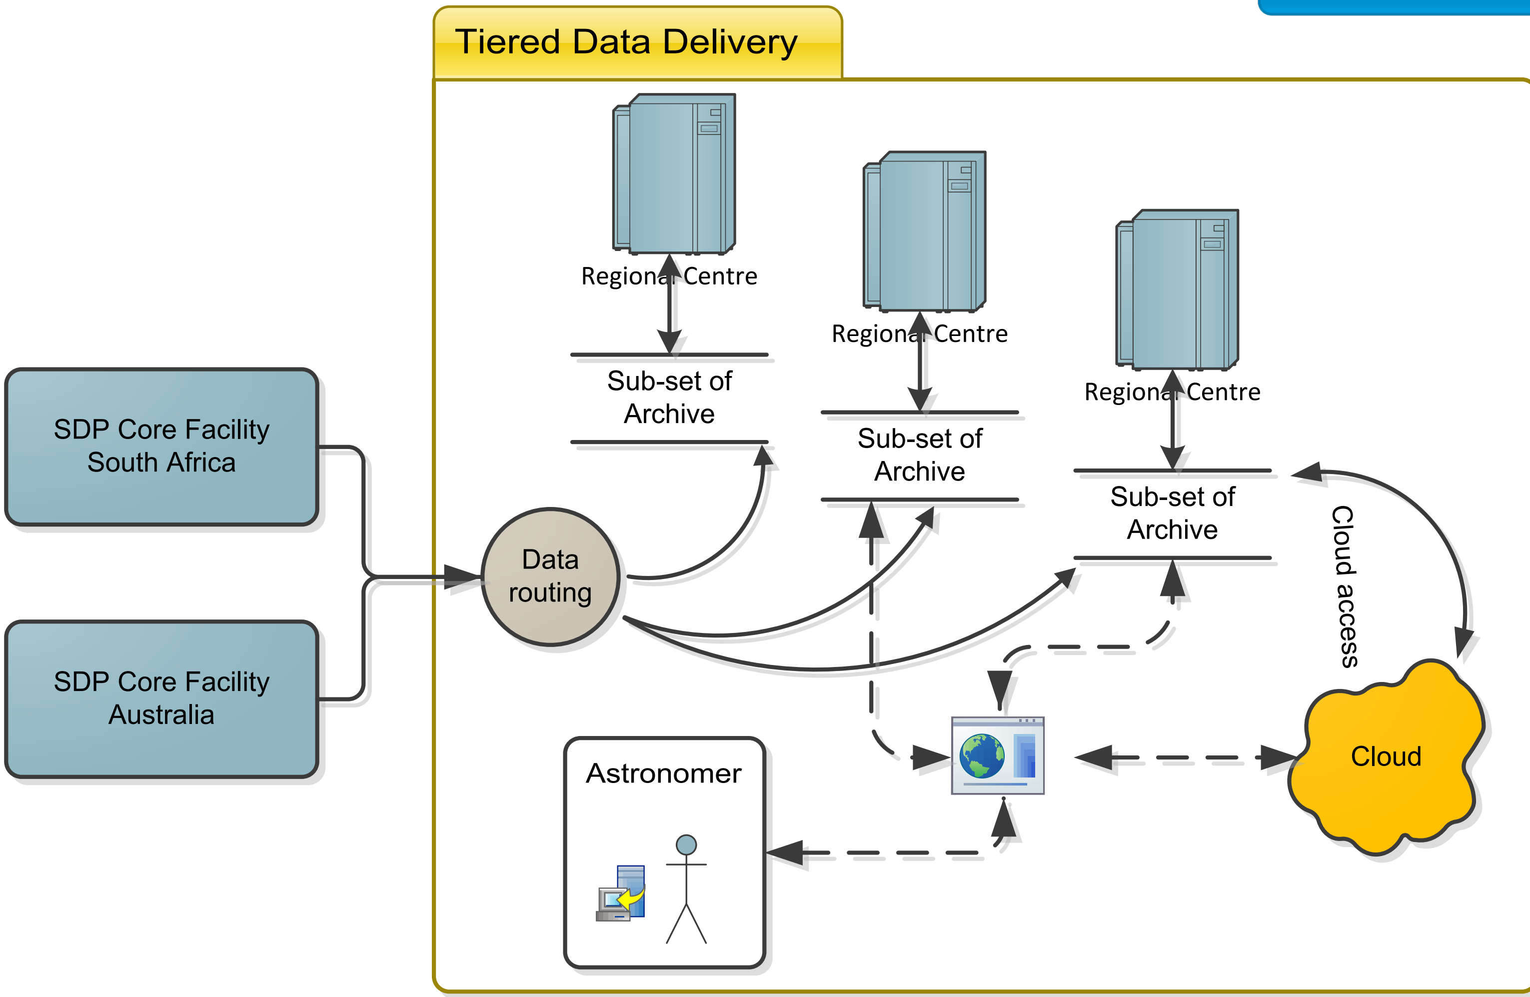

Regional Centres

- Last Week SKA Board Formally adopted concept of SKA Regional Centres
  - Data products from SKA up to 1PB/day
    - Still a major Big Data challenge
    - New tools and ideas needed
  - Provide access to SKA Data Products and Processing Environment
    - Typically manage 300 PB data/yr
    - Provide access to ~ 100 PFlop processing
  - Provide user support
  - Other roles
    - Continued input to SKA Observatory software support and development – Software Engineering Centre of Excellence
    - Development of framework for supporting Big Data challenge at Regional Centres
  - UK likely to be part of a European Centre using national infrastructures to provide the physical layer

Scope of the SDP

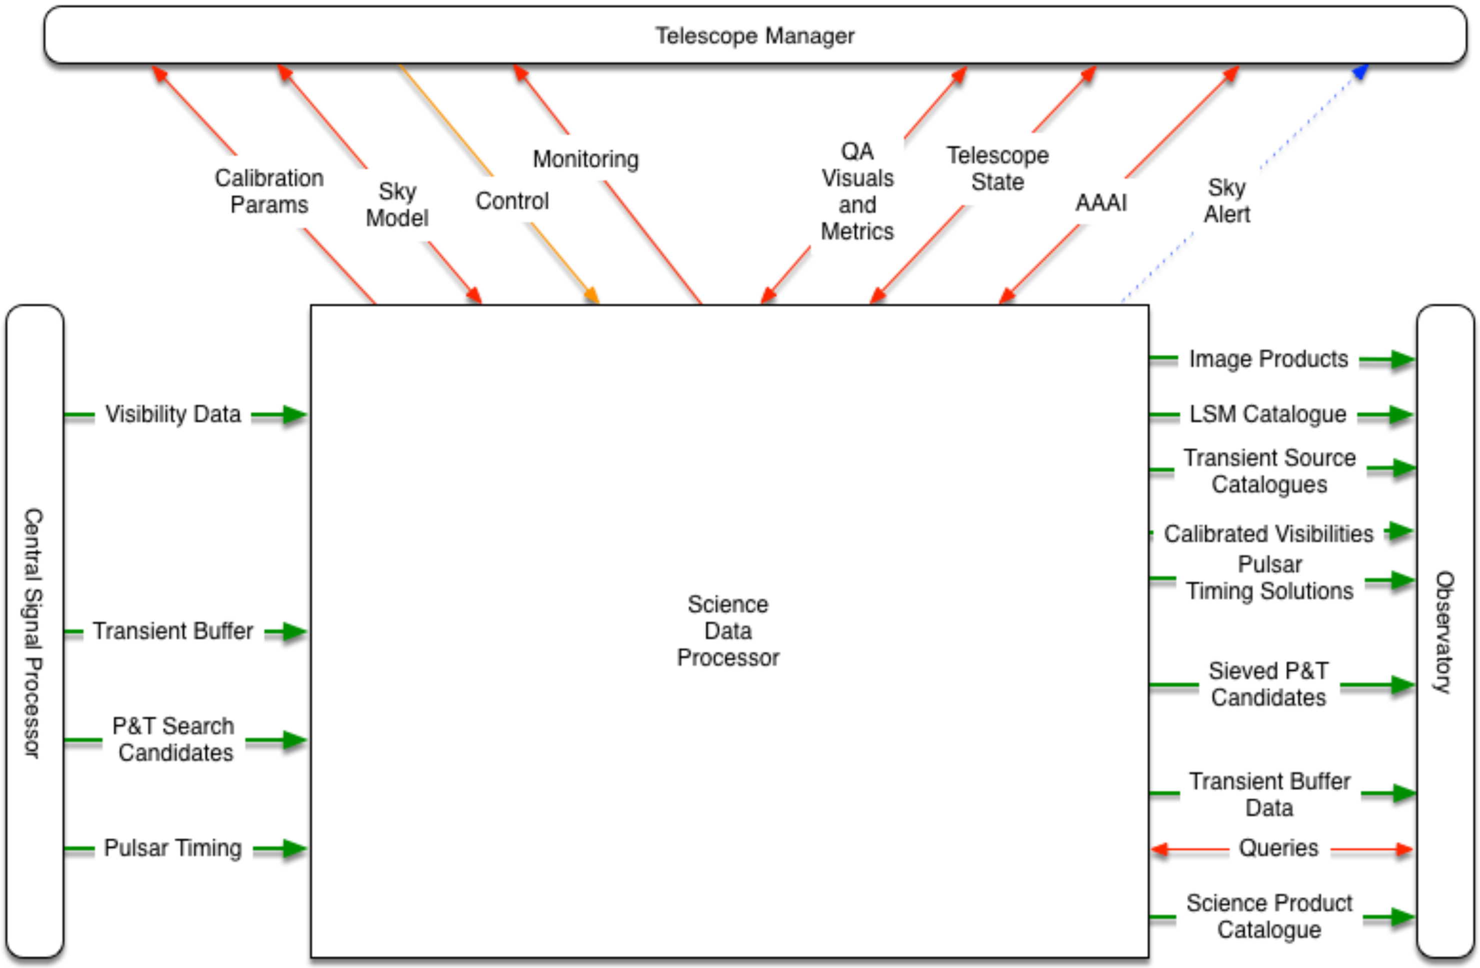

The SDP System

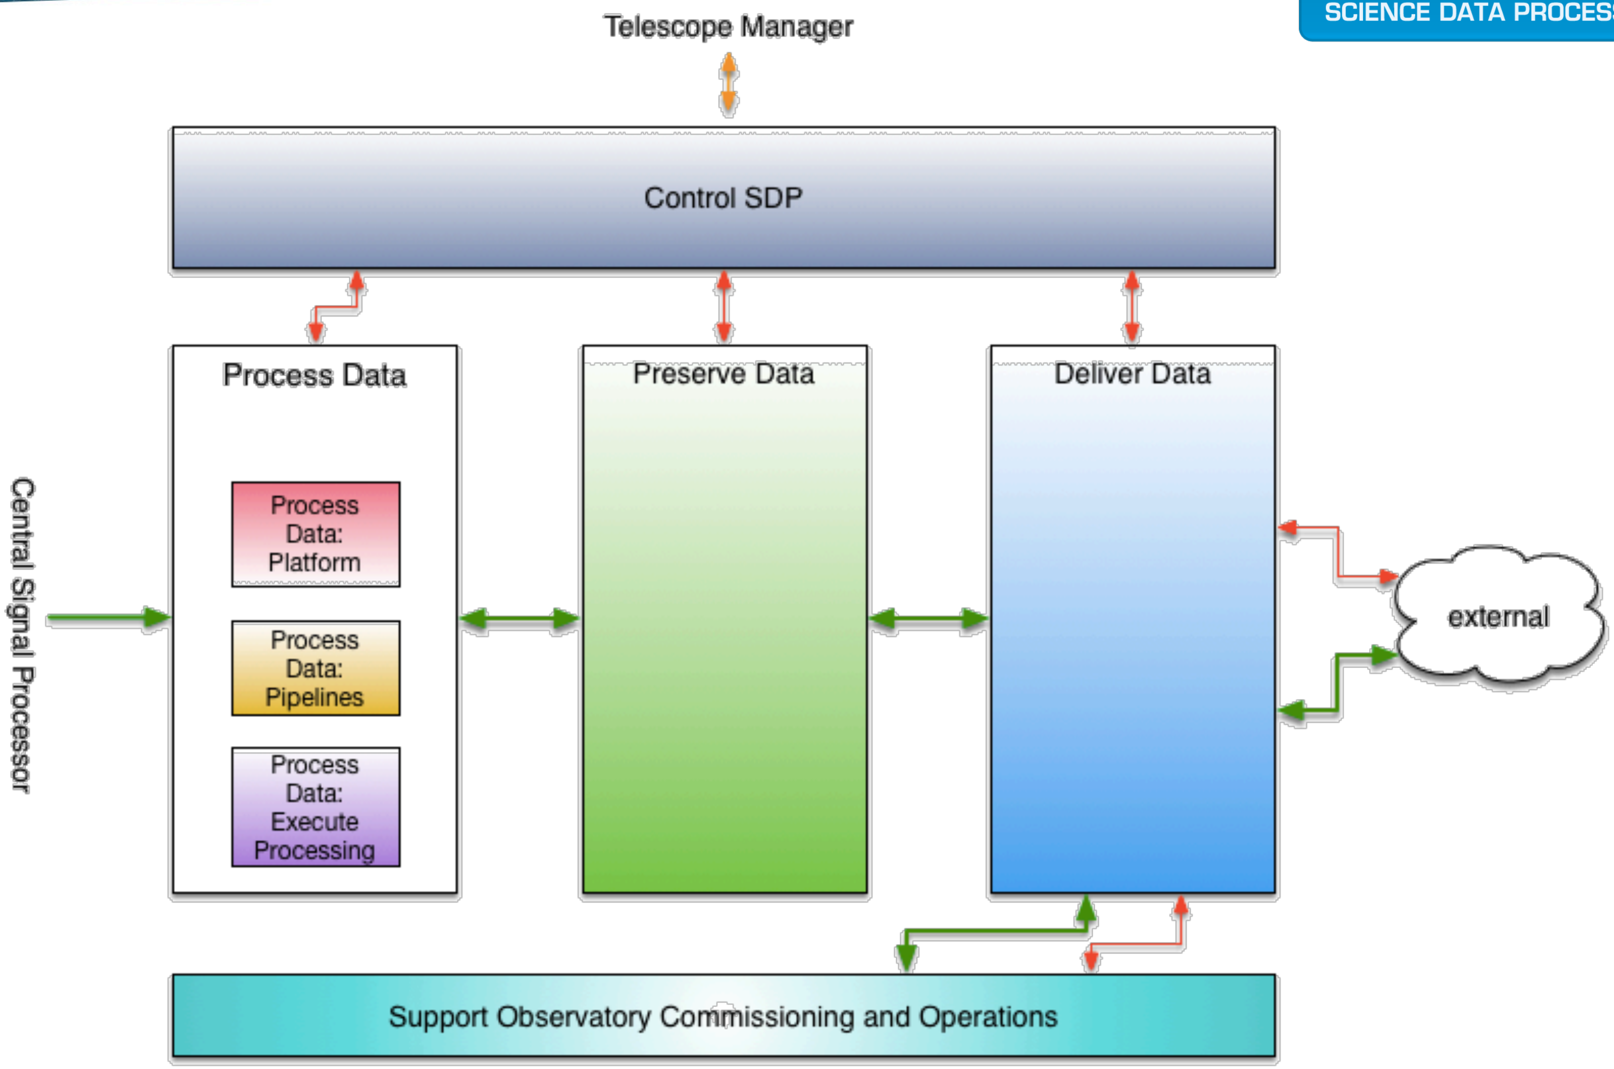

Architectural Approach

- Main principles
  - Ensure scalability (downwards mostly)
  - Ensure affordability
  - Ensure Maintainability
  - Support current state-of-the-art algorithms
- Exploit data parallelism, frequency & other dimensions
  - We have only two fundamental/bulk data structures
  - Raster grids and key-value-value stream records [e.g. u,v,w, -> visibility]
- Emphasis is on the framework to manage the throughput
  - Hardware platform will be replaced on a short duty cycle c.f. any HPC facility
  - Algorithms and workflow will evolve as we learn about telescopes

Approach: Co-design of software and physical layer architectures

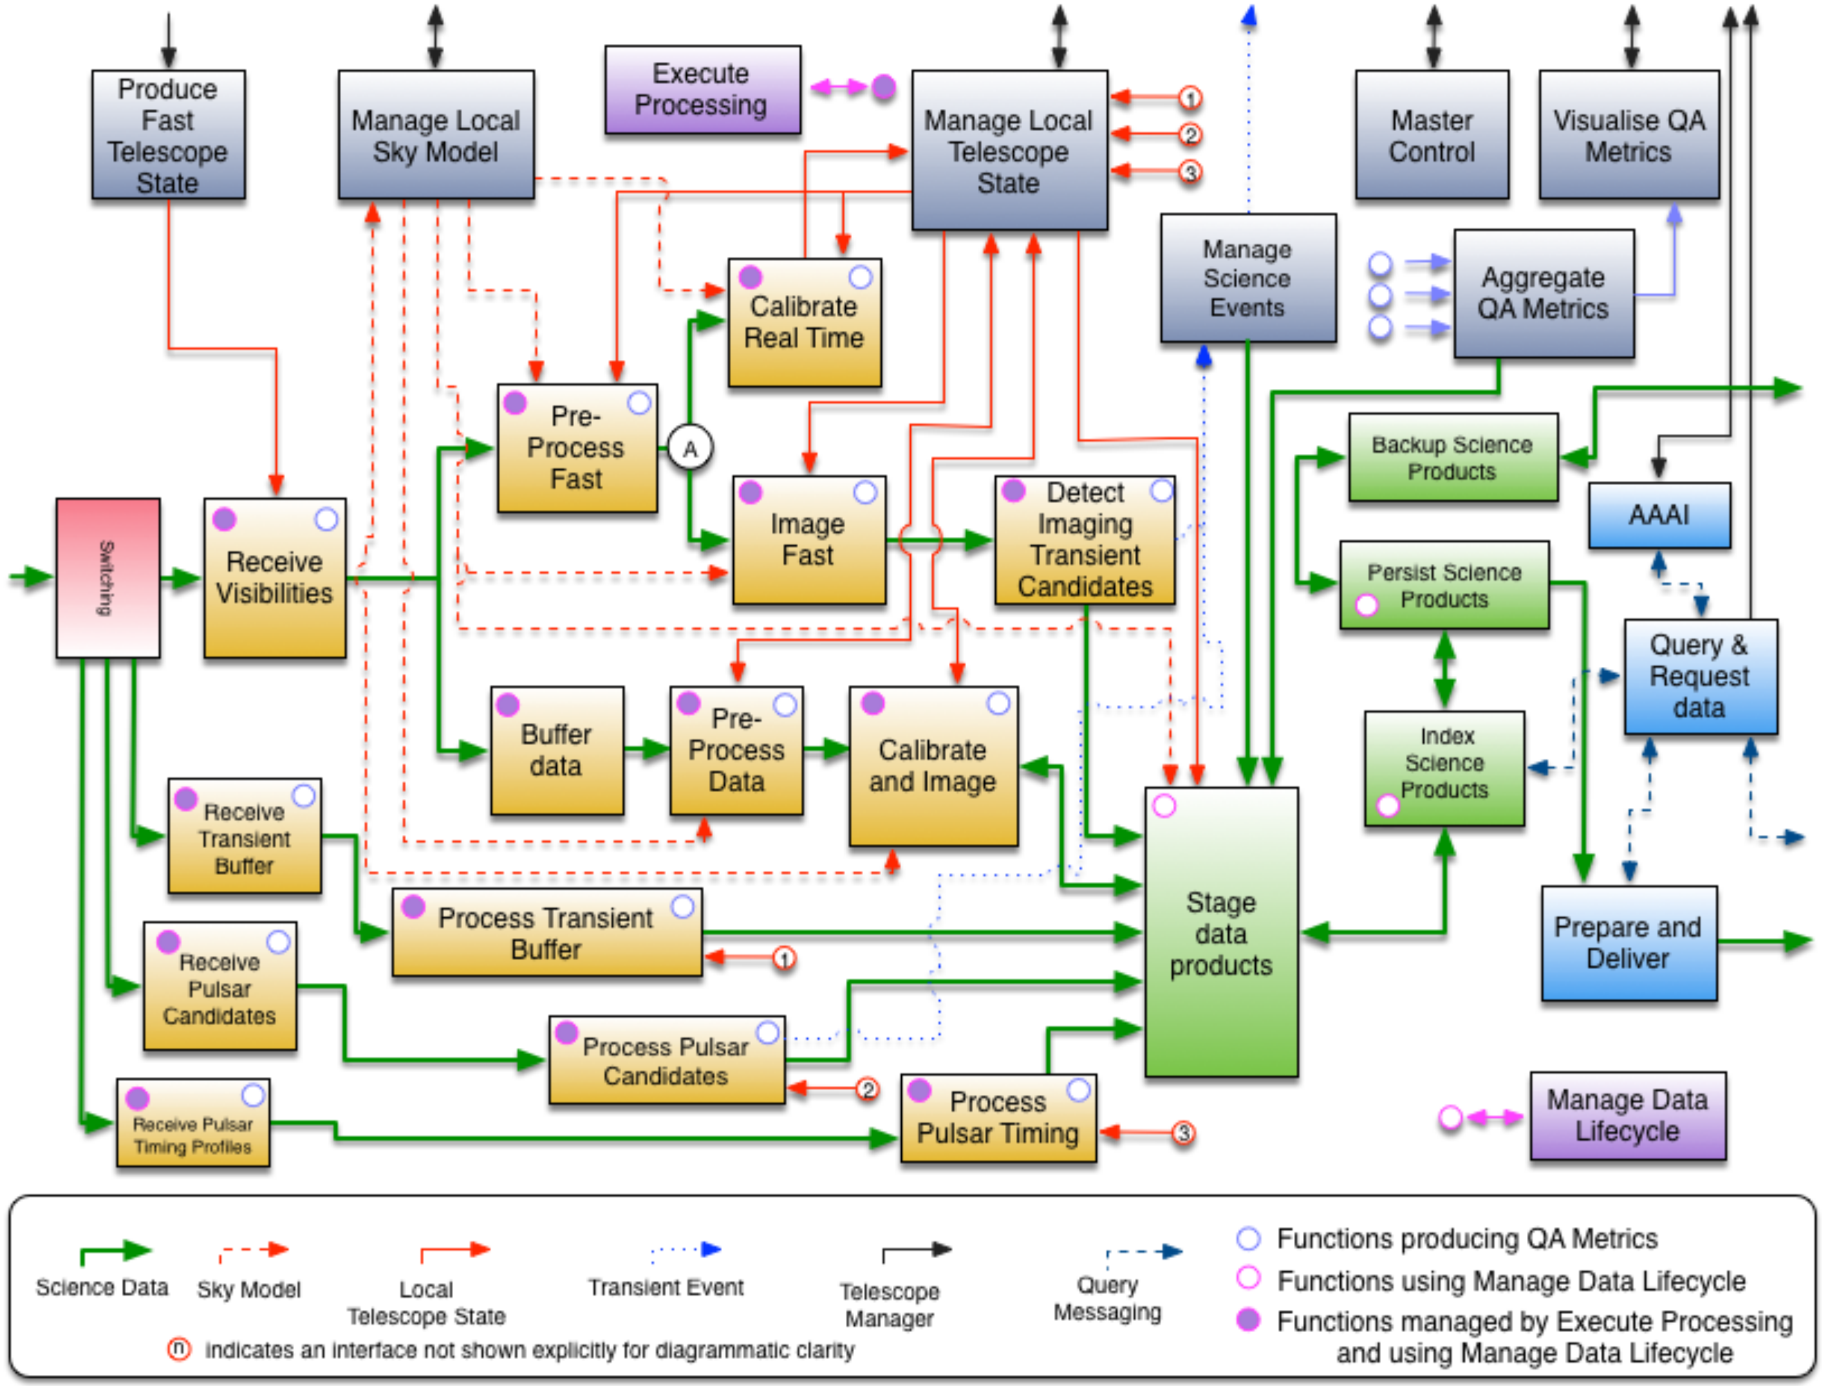

Produce Fast Telescope State
Manage Local Sky Model
Execute Processing
Manage Local Telescope State
Master Control
Visualise QA Metrics
Manage Science Events
Aggregate QA Metrics
Calibrate Real Time
Pre-Process Fast
A
Image Fast
Detect Imaging Transient Candidates
Backup Science Products
AAAI
Persist Science Products
Query & Request data
Index Science Products
Switching
Receive Visibilities
Buffer data
Pre-Process Data
Calibrate and Image
Stage data products
Receive Transient Buffer
Process Transient Buffer
Receive Pulsar Candidates
Process Pulsar Candidates
Receive Pulsar Timing Profiles
Process Pulsar Timing
Prepare and Deliver
Manage Data Lifecycle
1
2
3
Science Data
Sky Model
Local Telescope State
Transient Event
Telescope Manager
Query Messaging
Functions producing QA Metrics
Functions using Manage Data Lifecycle
Functions managed by Execute Processing and using Manage Data Lifecycle
n indicates an interface not shown explicitly for diagrammatic clarity

How do we get performance and manage data volume?

SCIENCE DATA PROCESSOR

Approach: Build on BigData Concepts

"data driven" → graph-based processing approach receiving a lot of attention

Inspired by Hadoop but for our complex data flow

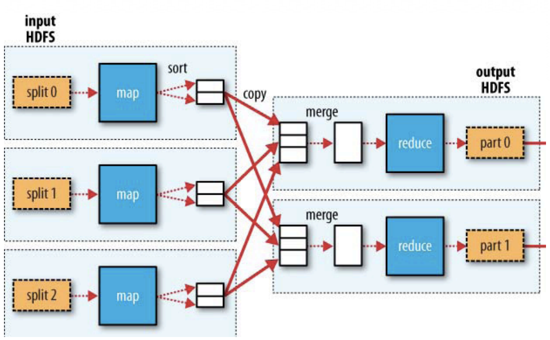

Hadoop

Graph-based approach

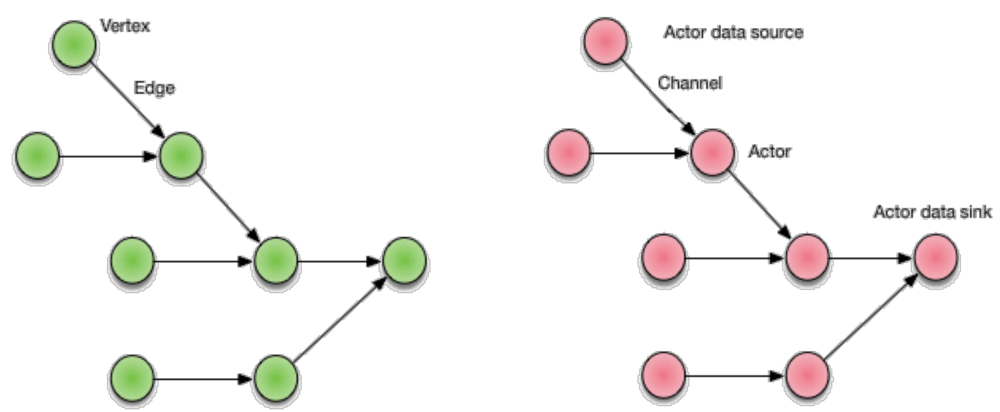

UNIVERSITY OF CAMBRIDGE

How do we get performance and manage data volume?

Key concepts:

- Exploit intrinsic data without strong coupling to the hardware Achieve suitable efficiency and in particular to load-balance the system;
- Separate domain-specific functionality from framework to give performance
- Represent pipeline as a graph showing data dependencies between components
- Components are the “tasks”, but now explicitly specify all of their required inputs and outputs and their execution is driven by the availability of data
  - This explicitly limits messages across the whole pipeline
  - Minimise data movement through the system by analysing the data dependencies and determining where to run processing
- Load balancing and scalability by task-based approach
- Fault tolerant
  - restarting processing based on data dependencies and reallocation of work

Data Driven Architecture

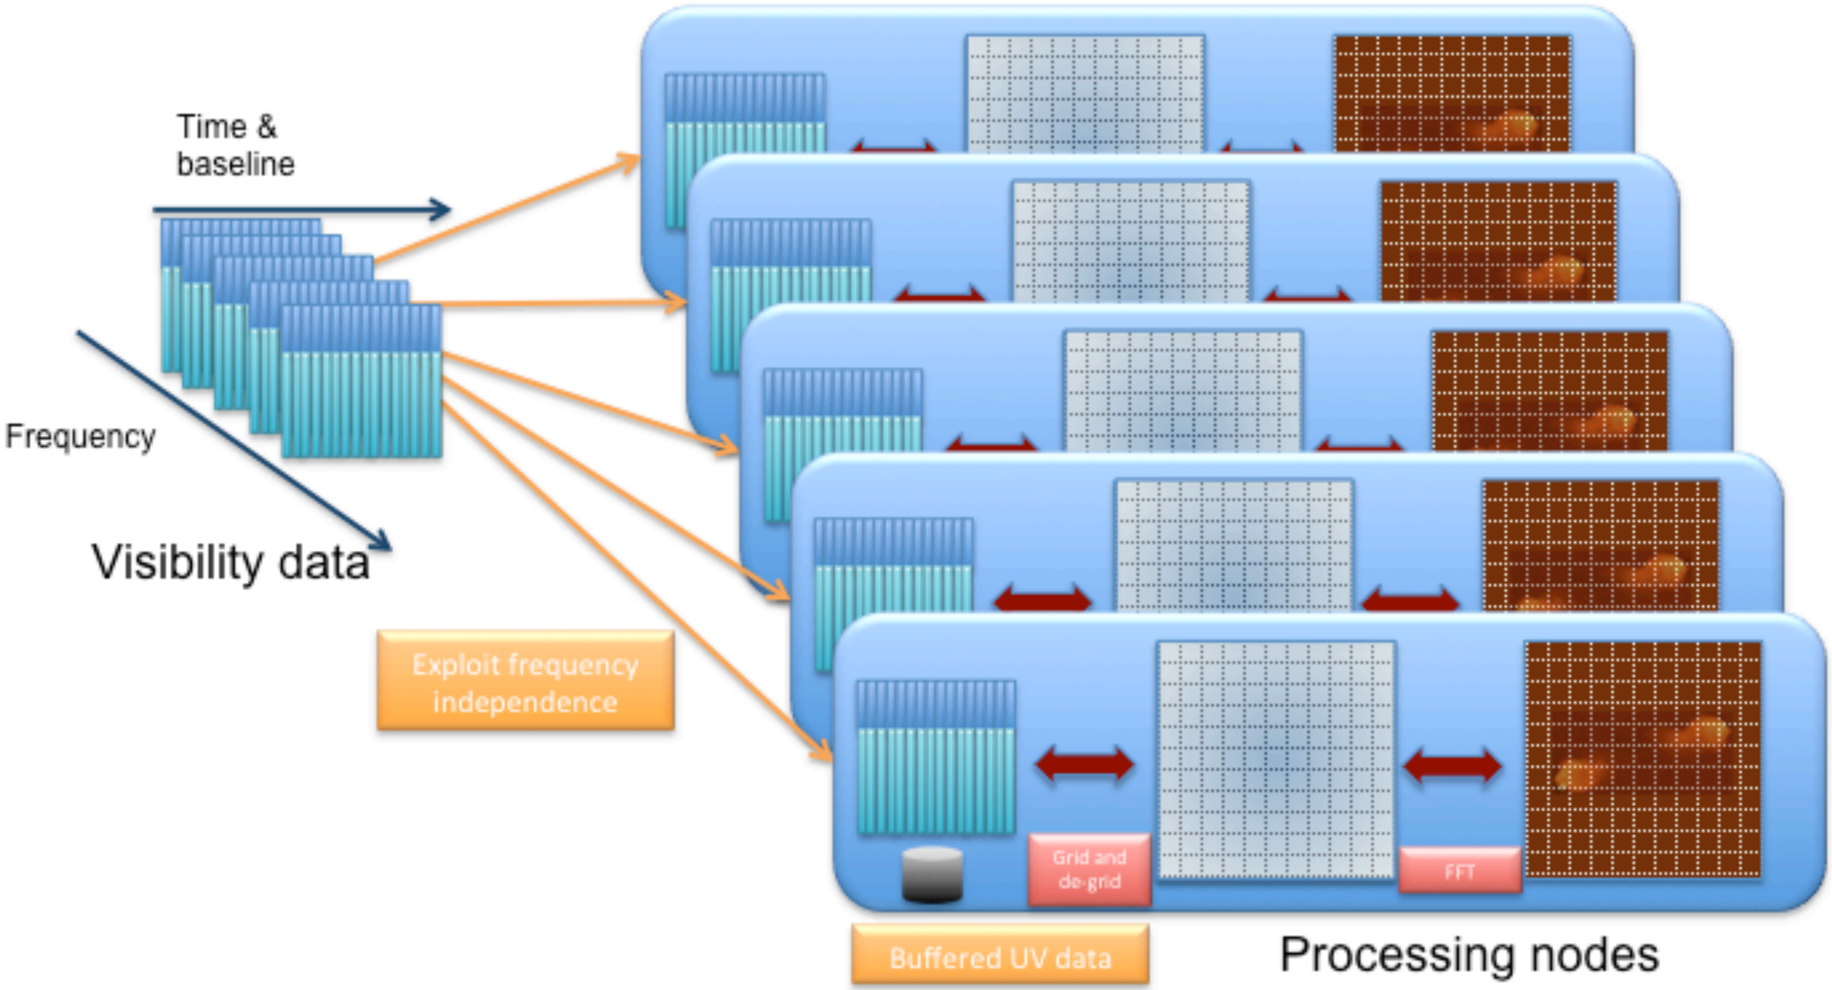

UNIVERSITY OF CAMBRIDGE

Data Driven Architecture

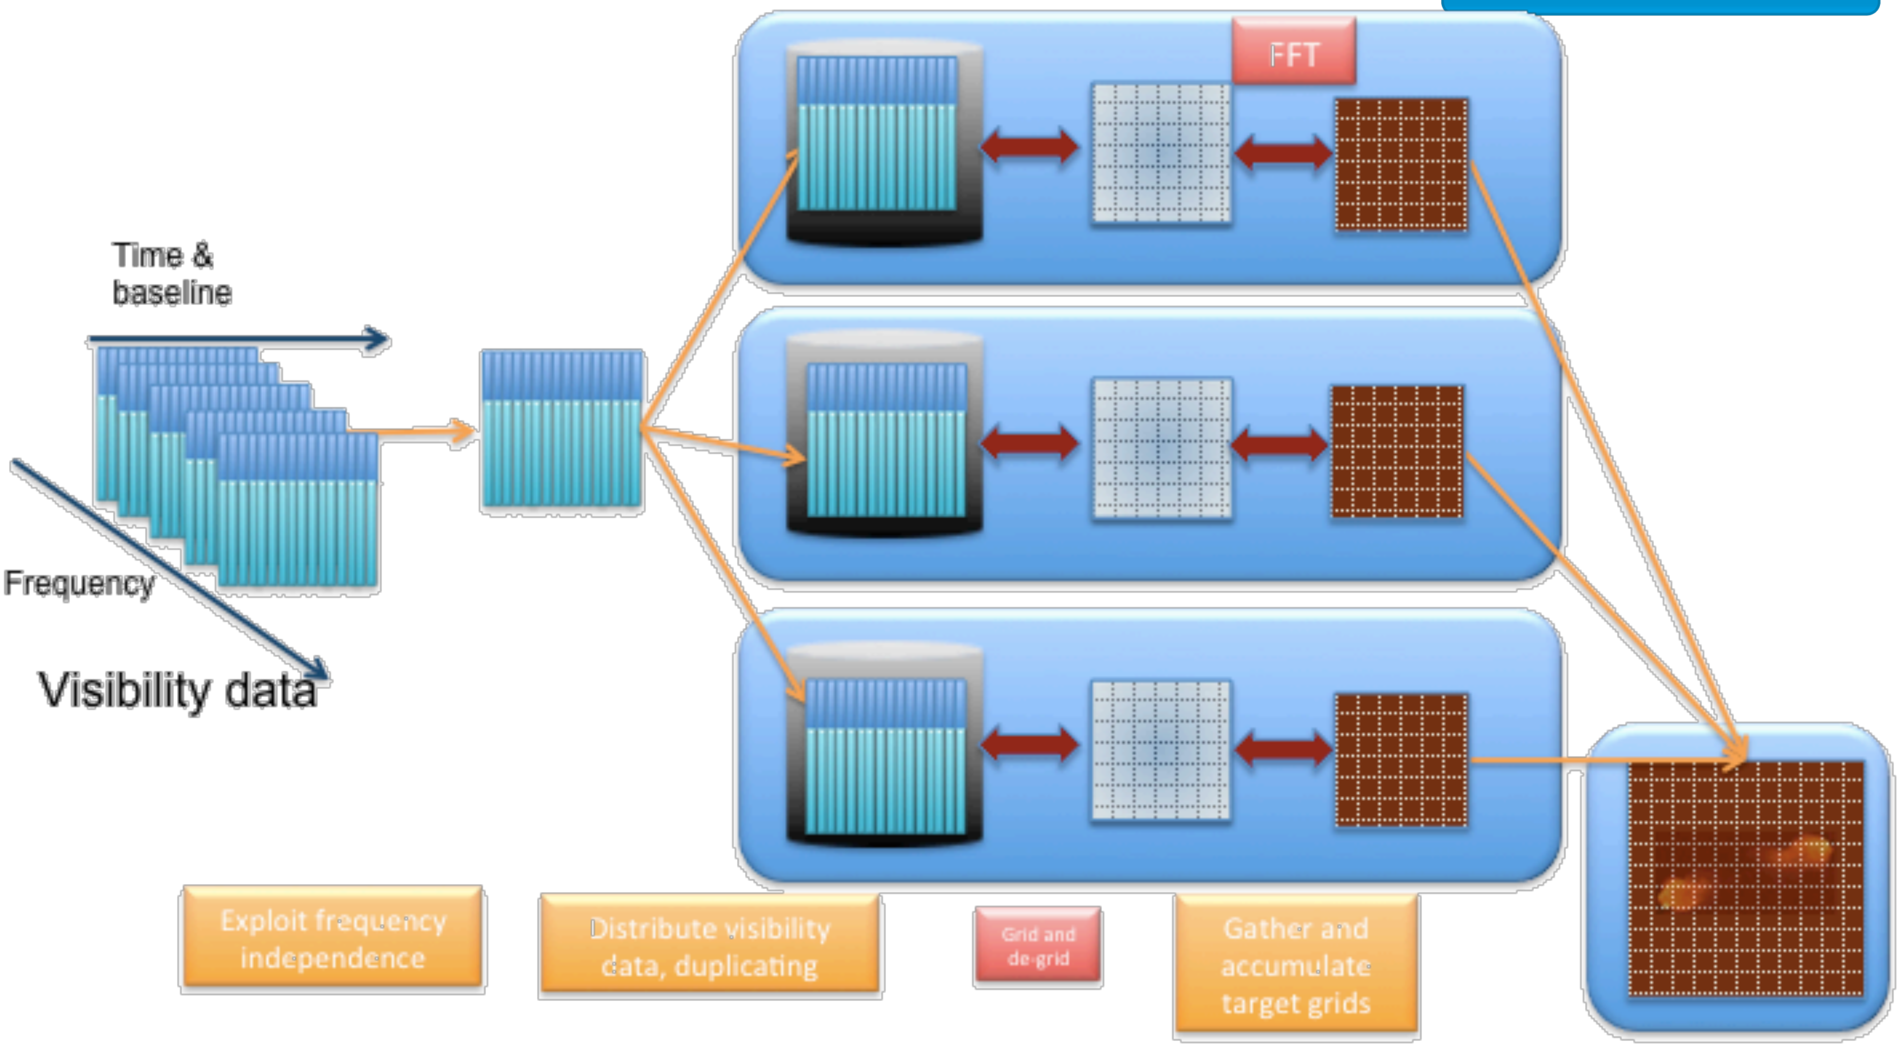

- Smaller FFT size at cost of data duplication

Pipelines as Graphs

Simplest example: Pre-processing data

- Split data by Beam/sub-array then time or frequency
- Load is split over many processors

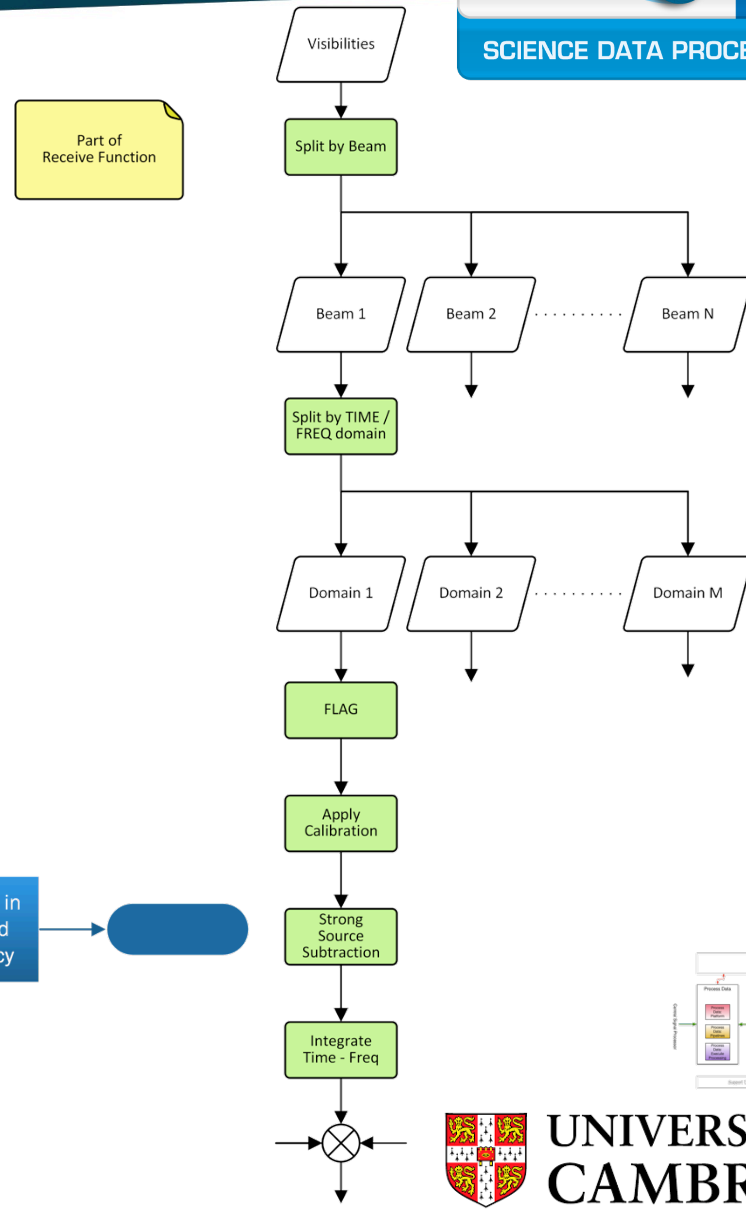

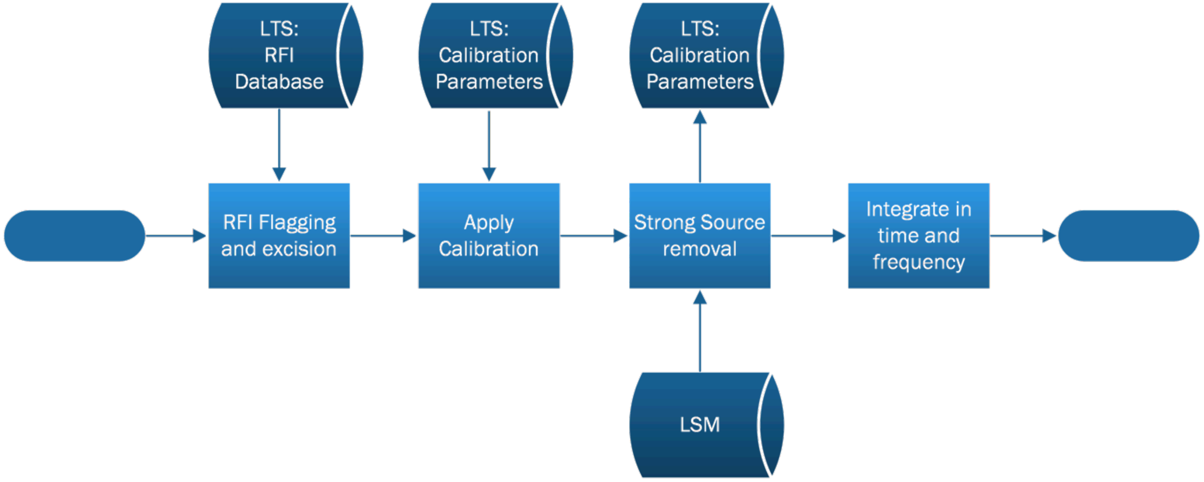

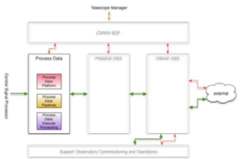

SCIENCE DATA PROCESSOR

UNIVERSITY OF CAMBRIDGE

SDP Data Flow Approach: Next Generation Data Science Engine?

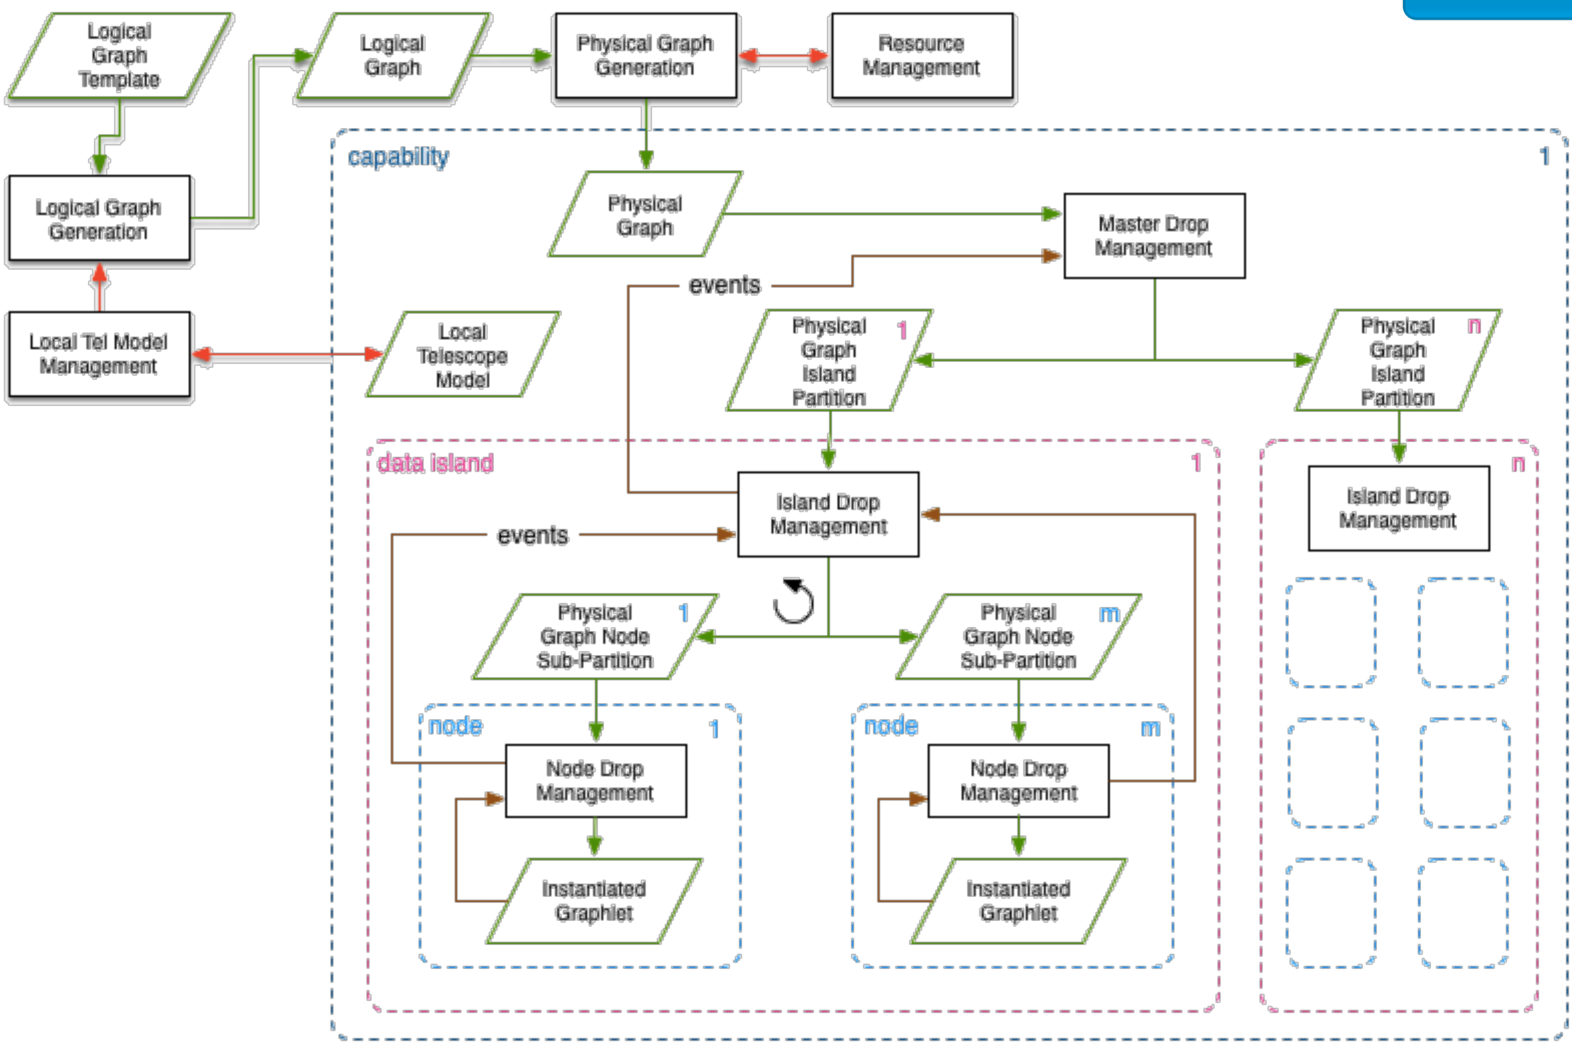

SKA SQUARE KILOMETRE ARRAY
SDP
SCIENCE DATA PROCESSOR

UNIVERSITY OF CAMBRIDGE

Managing data and separating functionality

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

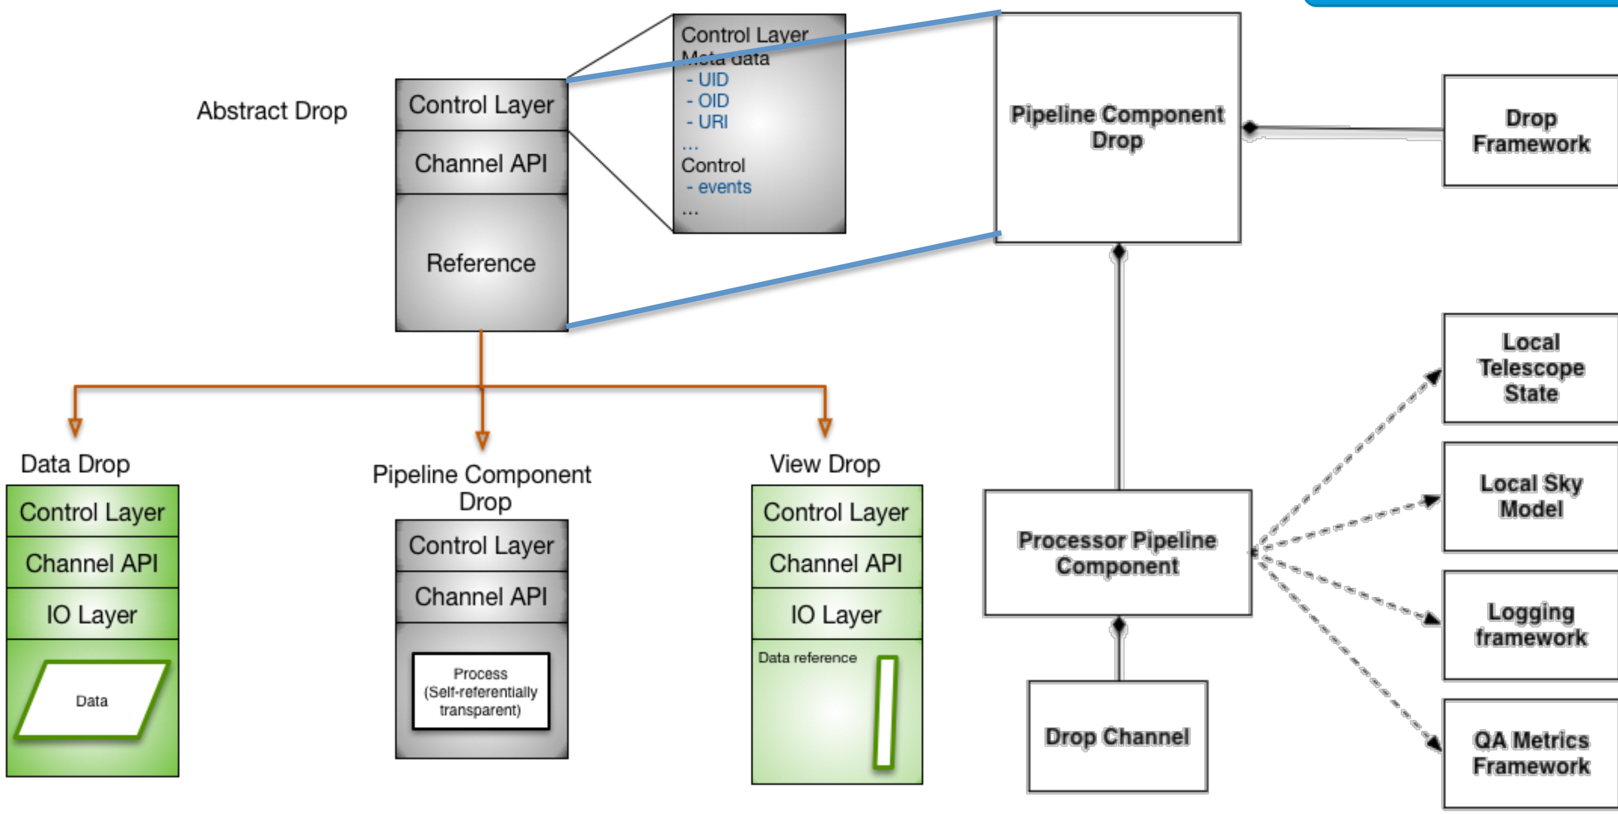

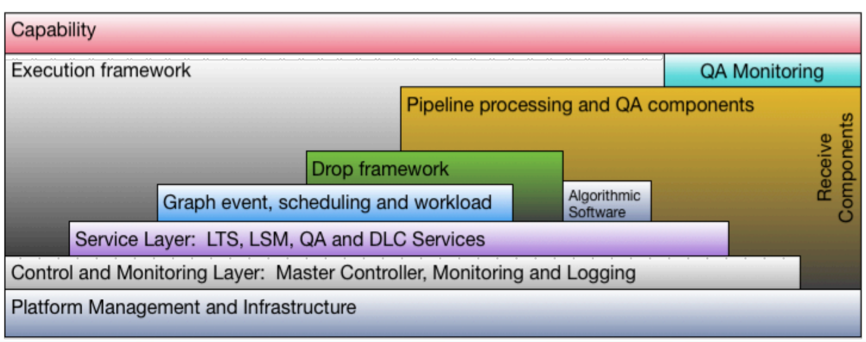

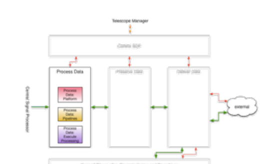

UNIVERSITY OF CAMBRIDGE

The graph in operation

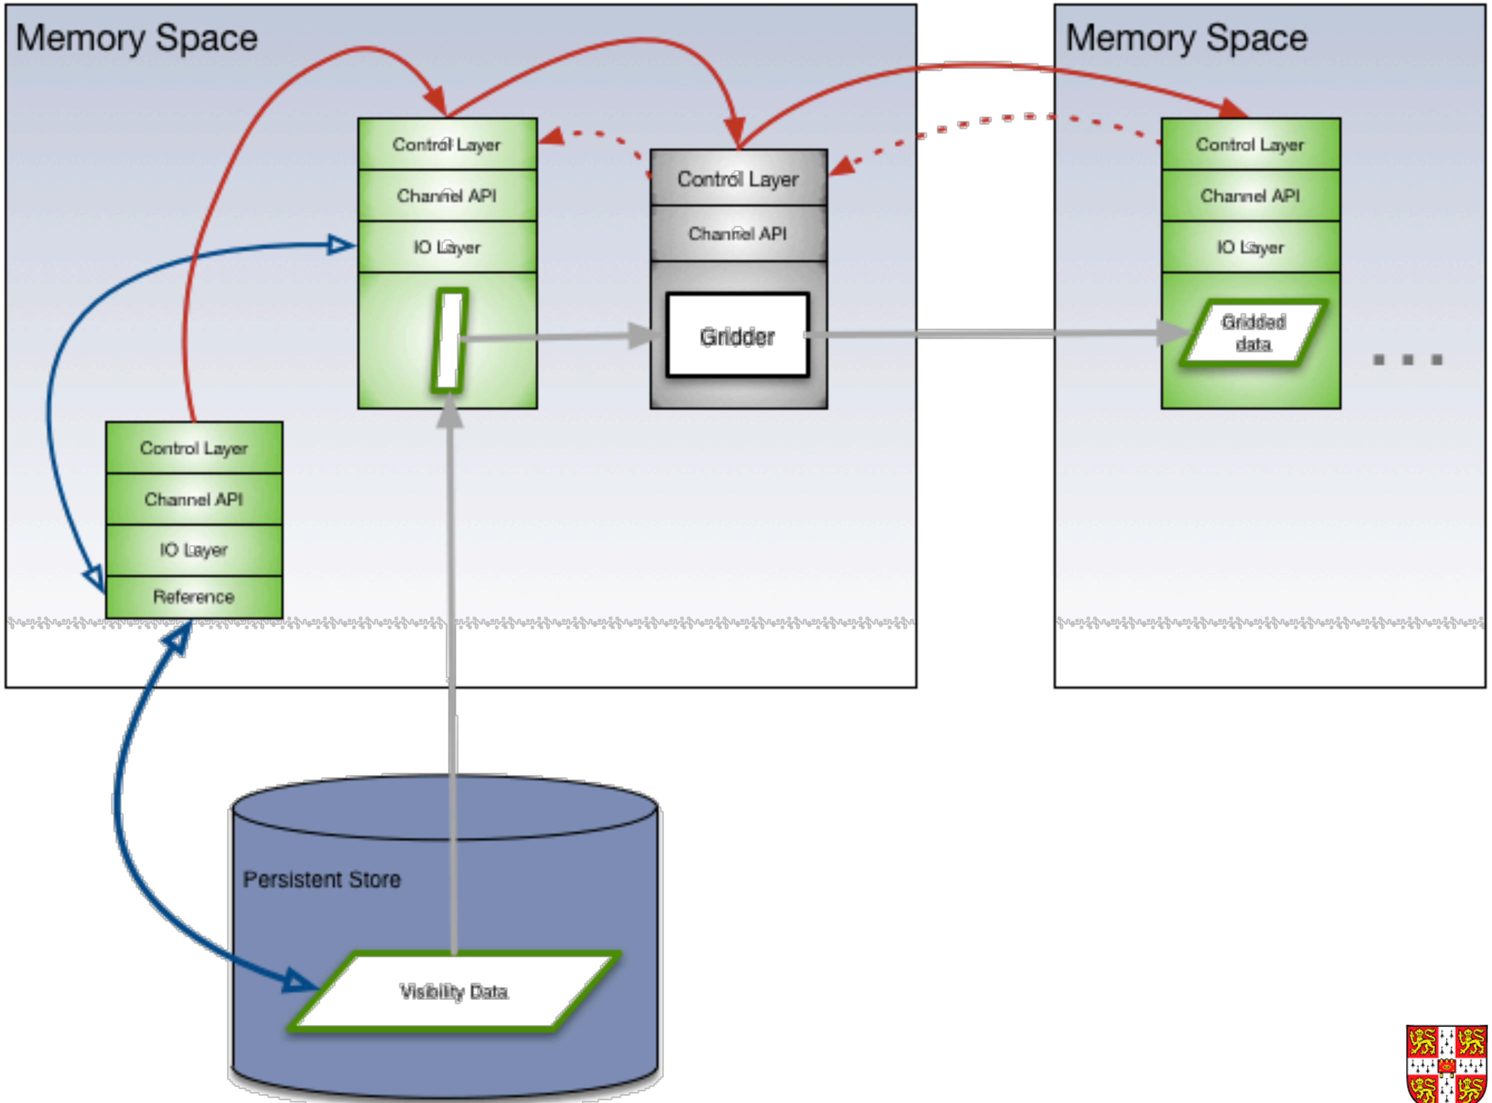

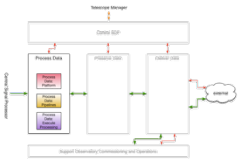

UNIVERSITY OF CAMBRIDGE

Controlling SDP

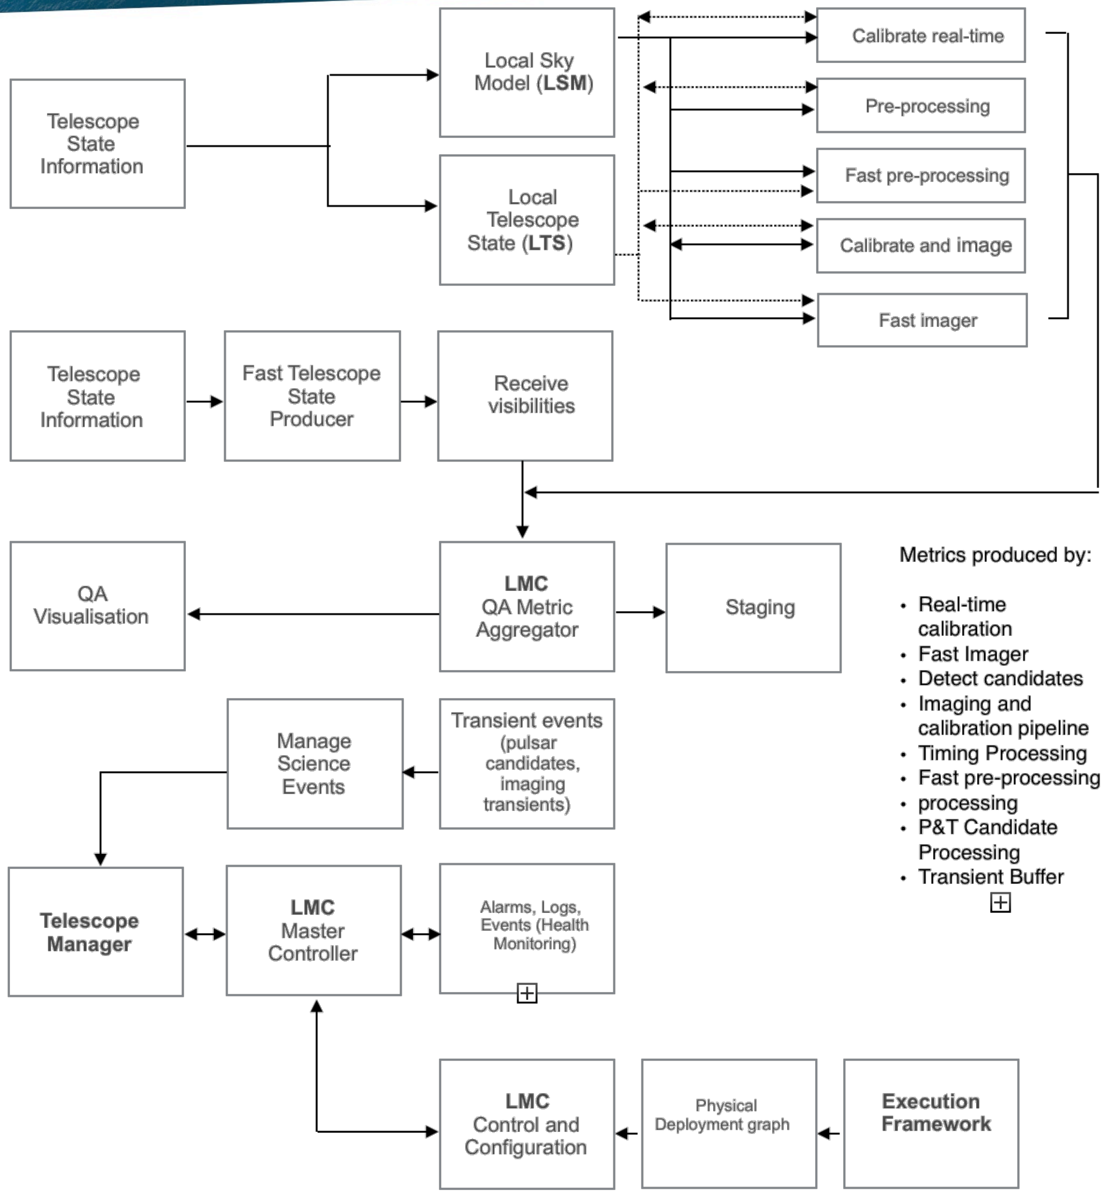

Metrics produced by:

- Real-time calibration
- Fast Imager
- Detect candidates
- Imaging and calibration pipeline
- Timing Processing
- Fast pre-processing
- processing
- P&T Candidate Processing
- Transient Buffer

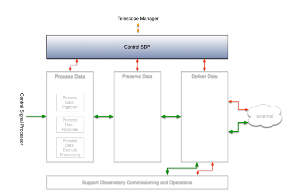

Important Services: Quality Assessment, Sky Models and Telescope State

SCIENCE DATA PROCESSOR

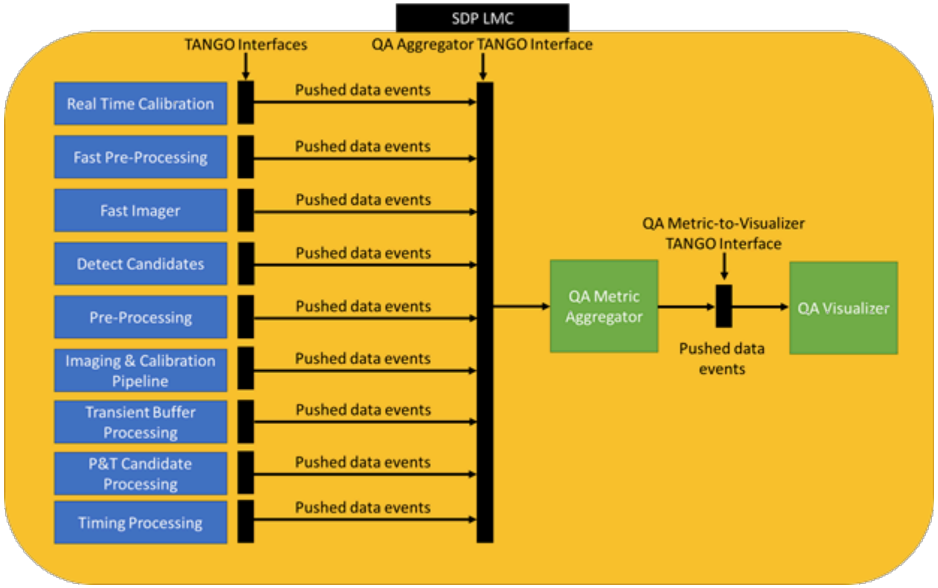

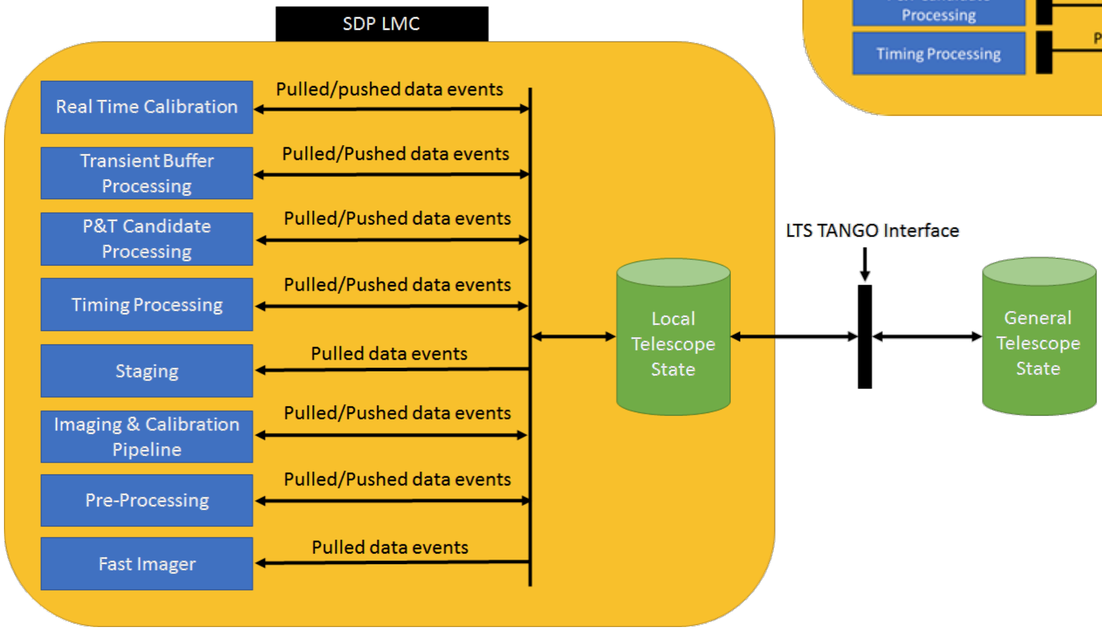

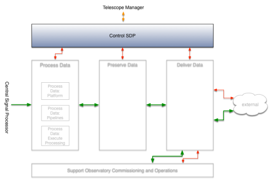

UNIVERSITY OF CAMBRIDGE

Delivering Data

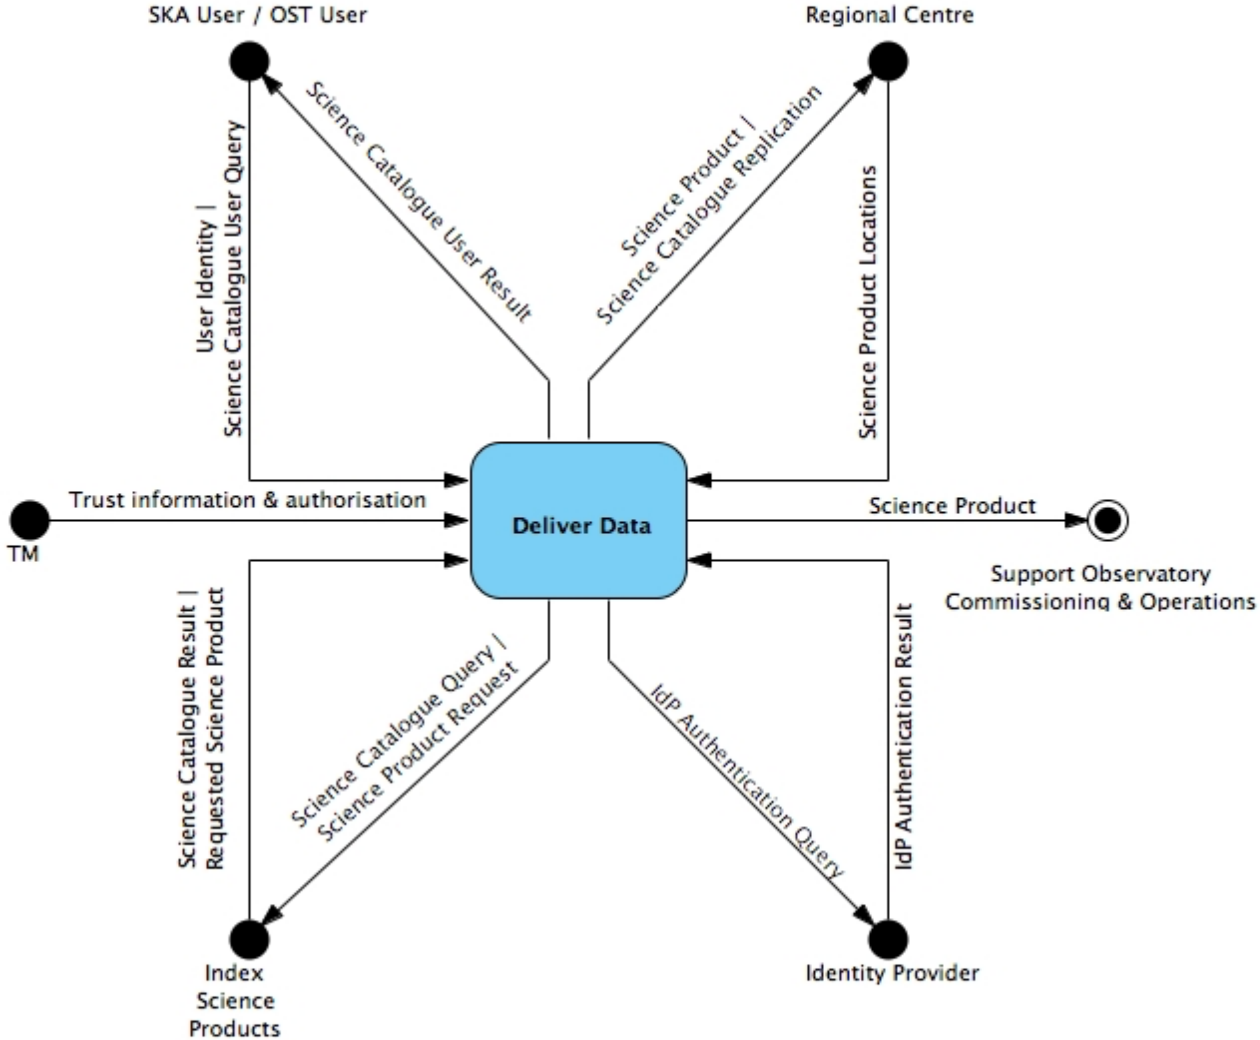

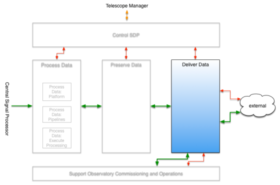

Hardware Platform

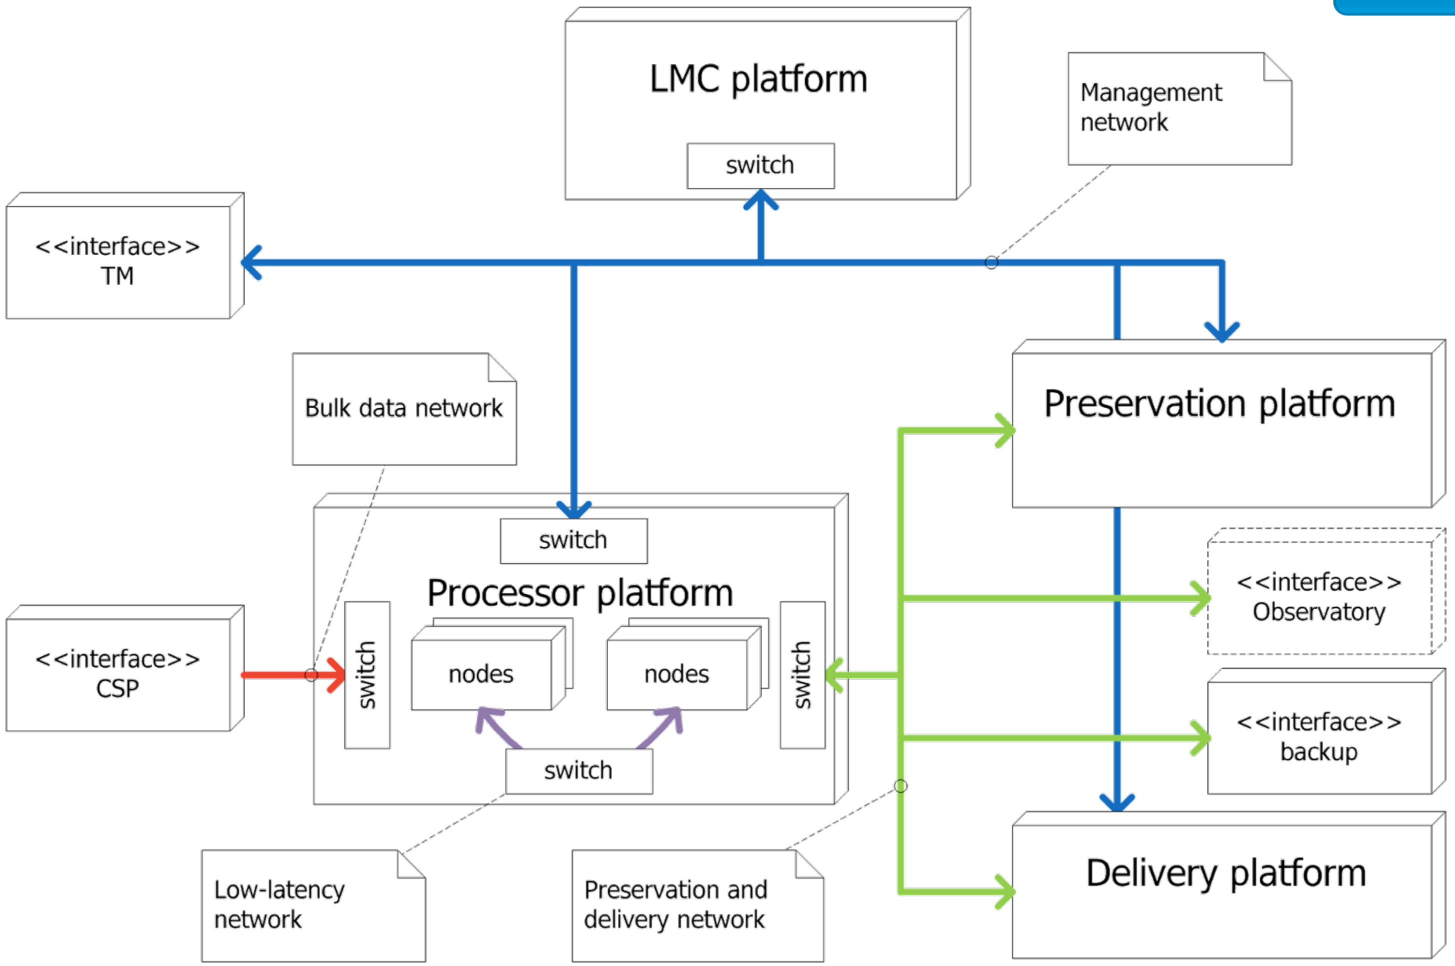

UNIVERSITY OF CAMBRIDGE

SCIENCE DATA PROCESSOR

Data rates and processing increase by FACTOR ~100 for SKA2

3-30 EBytes / year of fully processed data for SKA2

UNIVERSITY OF CAMBRIDGE

END

Complex Network

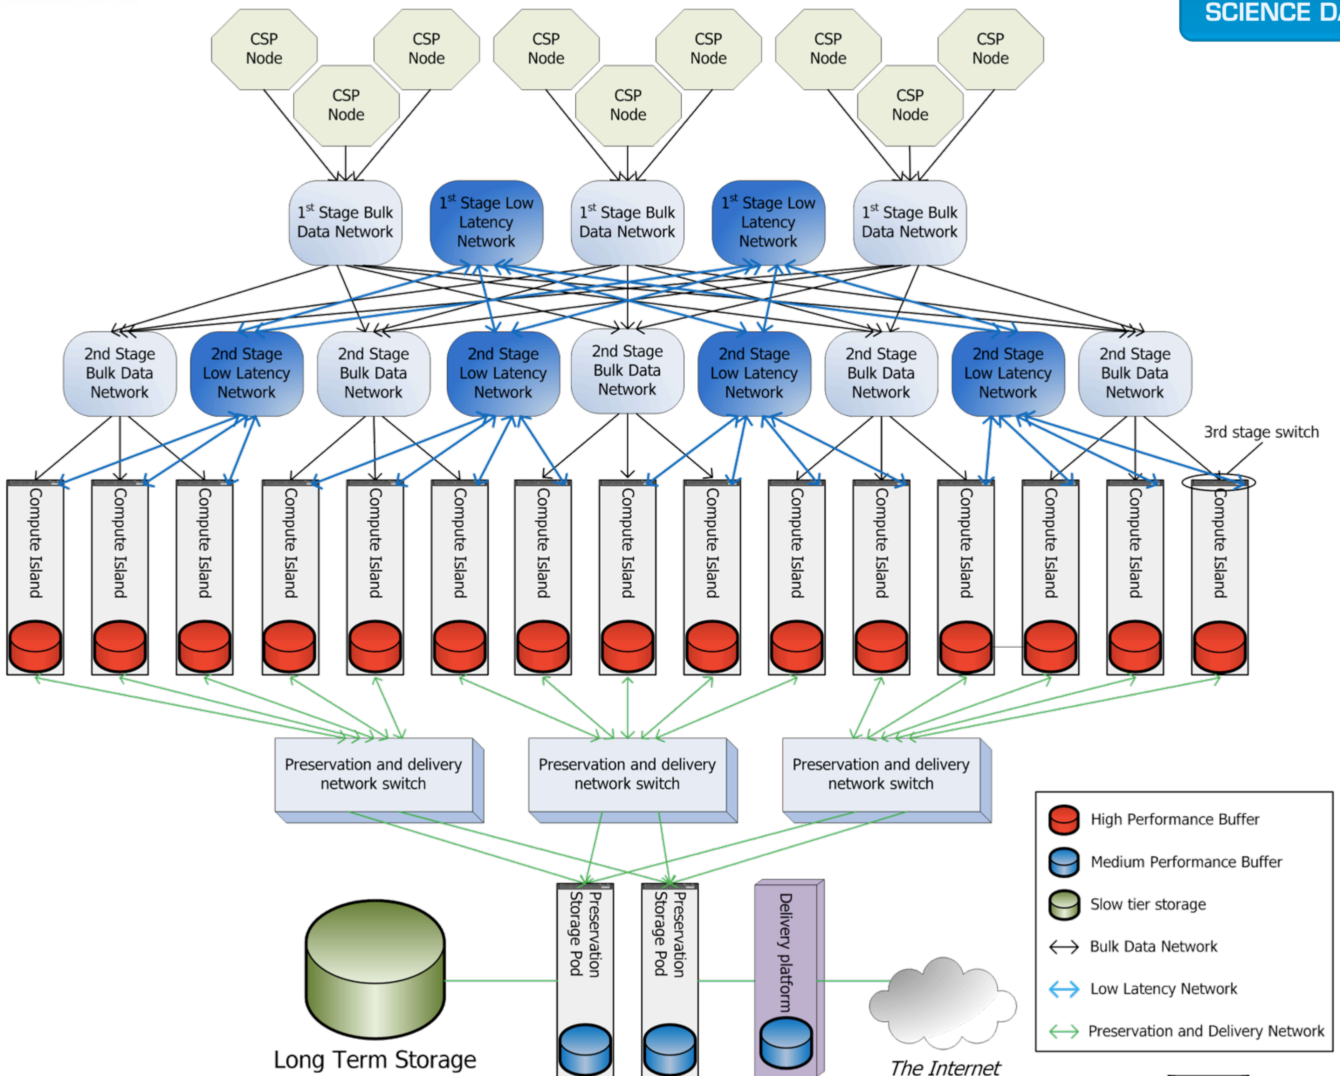

Data Driven Architecture

SKA SQUARE KILOMETRE ARRAY SDP
SCIENCE DATA PROCESSOR

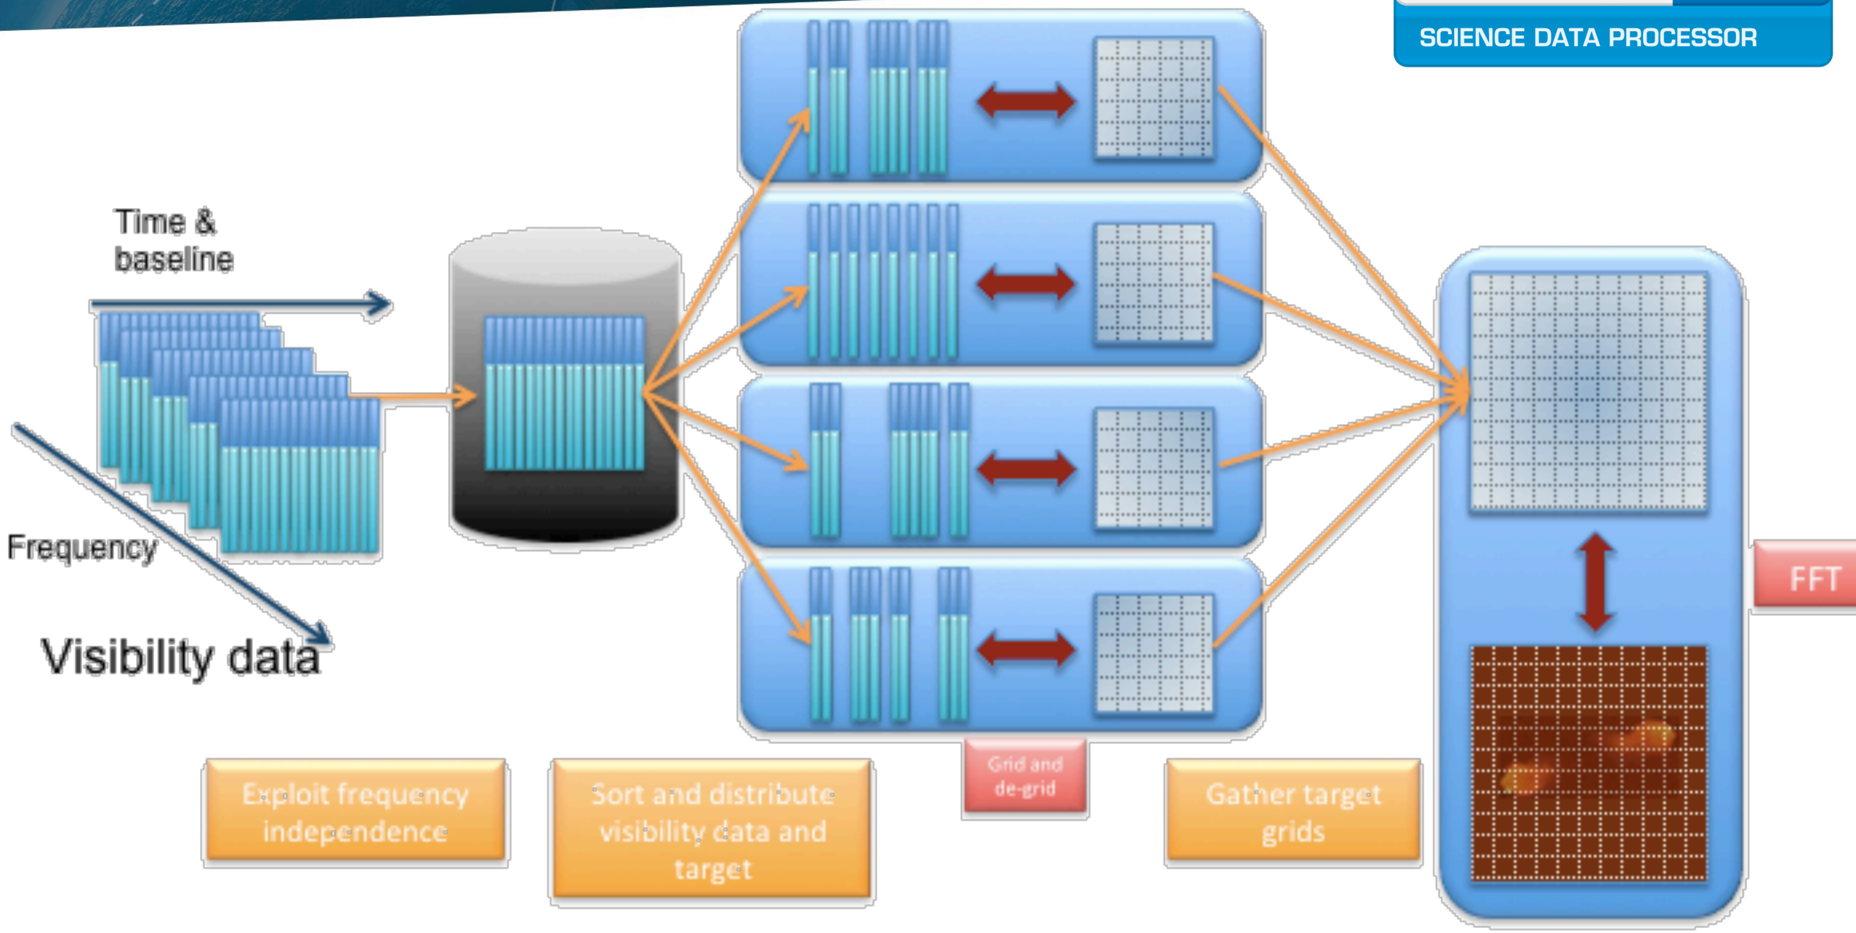

- Further data parallelism in spatial indexing (UVW-space)
- Use to balance memory bandwidth per node
- Some overlap regions on target grids needed

UNIVERSITY OF CAMBRIDGE

Data Driven Architecture

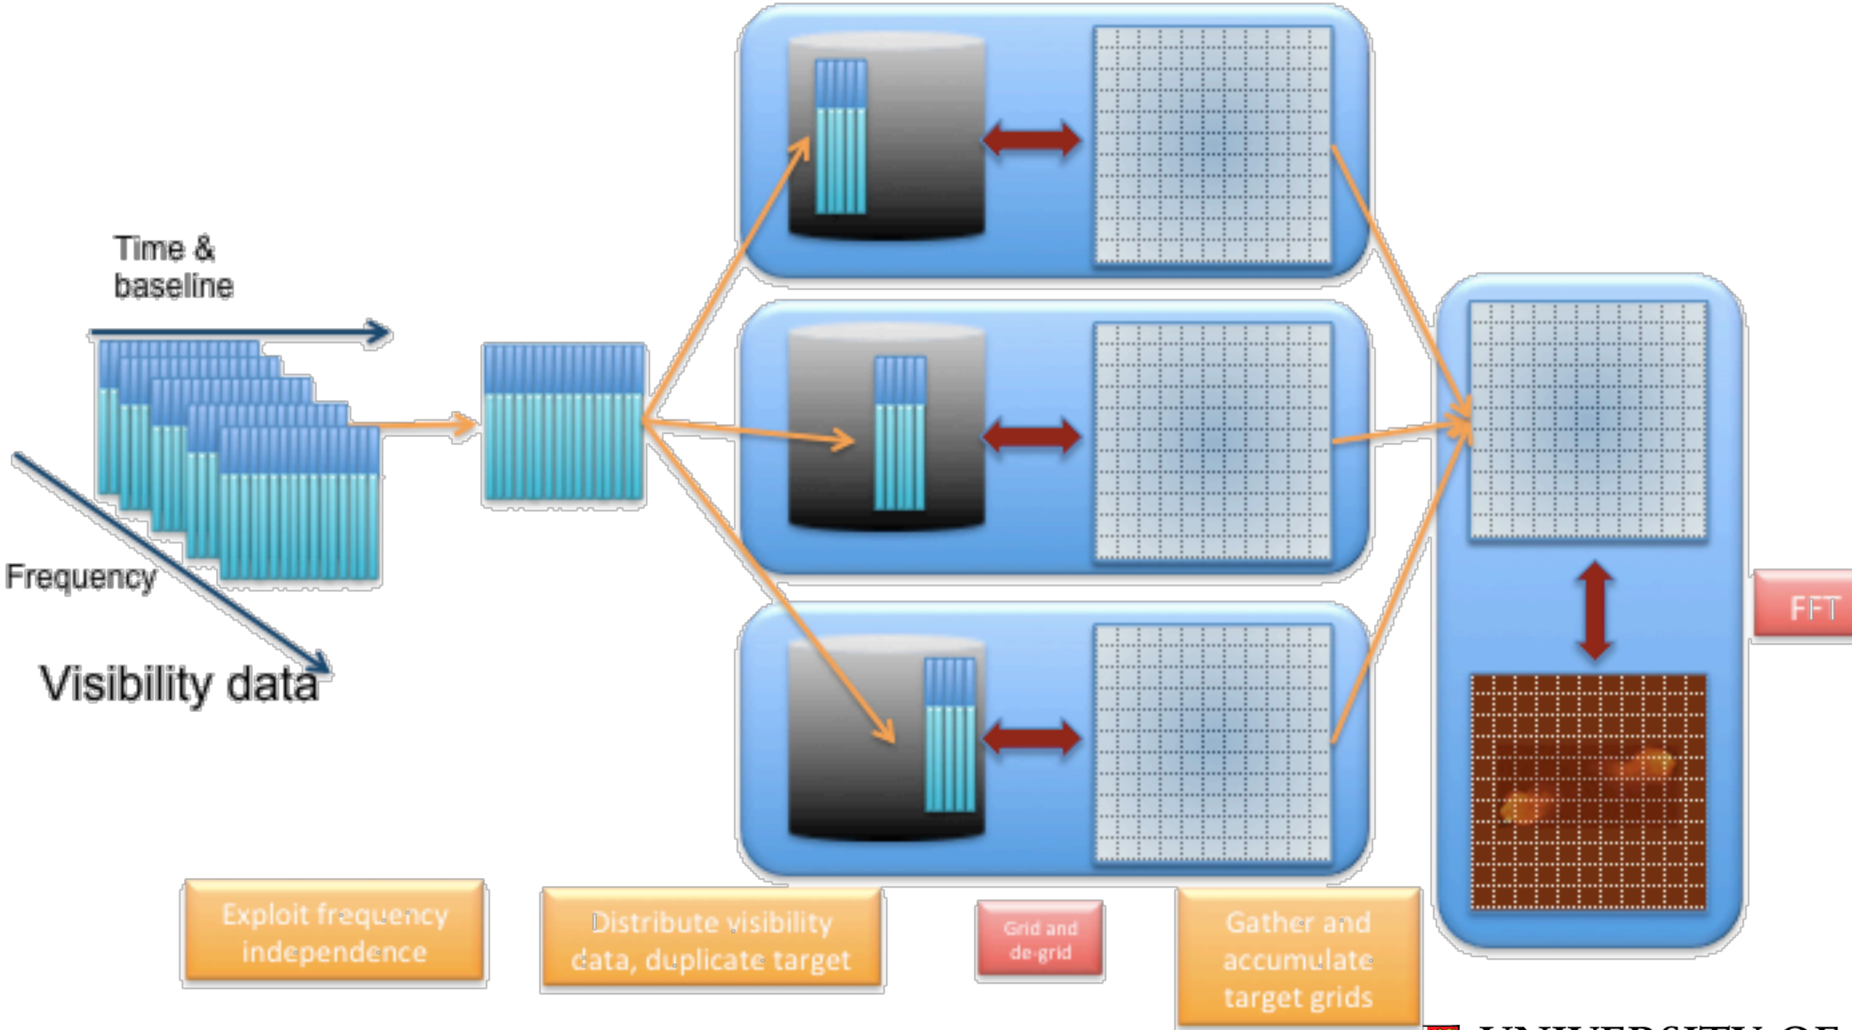

UNIVERSITY OF CAMBRIDGE