The path to the SKA: Big data challenges in radio astronomy

Dr. Ewan Barr
SKA CSP Project Scientist
Max Planck Institut für Radioastronomie (MPIfR)
Bonn

Overview

- Why bigger data?
- SKA overview
- Big data challenges for the SKA
- Precursors and technology demonstrators
  - MeerKAT S-band receiver / backend project
- Summary

Why bigger data?

- What do we want:
    - Higher sensitivity
    - Higher time resolution
    - Higher spatial resolution
    - Wider fields of view (FoV)
    - Higher efficiency

Bandwidth

Bigger dishes

Smaller dishes

Rawer data

Why bigger data?

- What do we want:
  - Higher sensitivity
  - Higher time resolution
  - Higher spatial resolution
  - Wider fields of view (FoV)
  - Higher efficiency

Bandwidth

Many small dishes:
Small D - Large N

Rawer data

Why bigger data?

Bandwidth → High data rate streams

Many small dishes: Small D - Large N → $O(N^2)$ processing

Rawer data → Large storage requirements

What happens when you take this to the extreme?


SKA
SQUARE KILOMETRE ARRAY
www.skatelescope.org
The Square Kilometre Array
SWINBURNE ASTRONOMY PRODUCTIONS

SKA Phase 1 Overview

SKA1 Low

SKA1 Mid

- Wide-FoV "software" telescopes
- Capable of performing multiple projects simultaneously
- International mega-science project (10+ member states)

SKA Organisation: 10 countries, more to join

Australia (DoI&S)
Canada (NRC-HIA)
China (MOST)
India (DAE)
Italy (INAF)
Netherlands (NWO)
New Zealand (MED)
South Africa (DST)
Sweden (Chalmers)
UK (STFC)

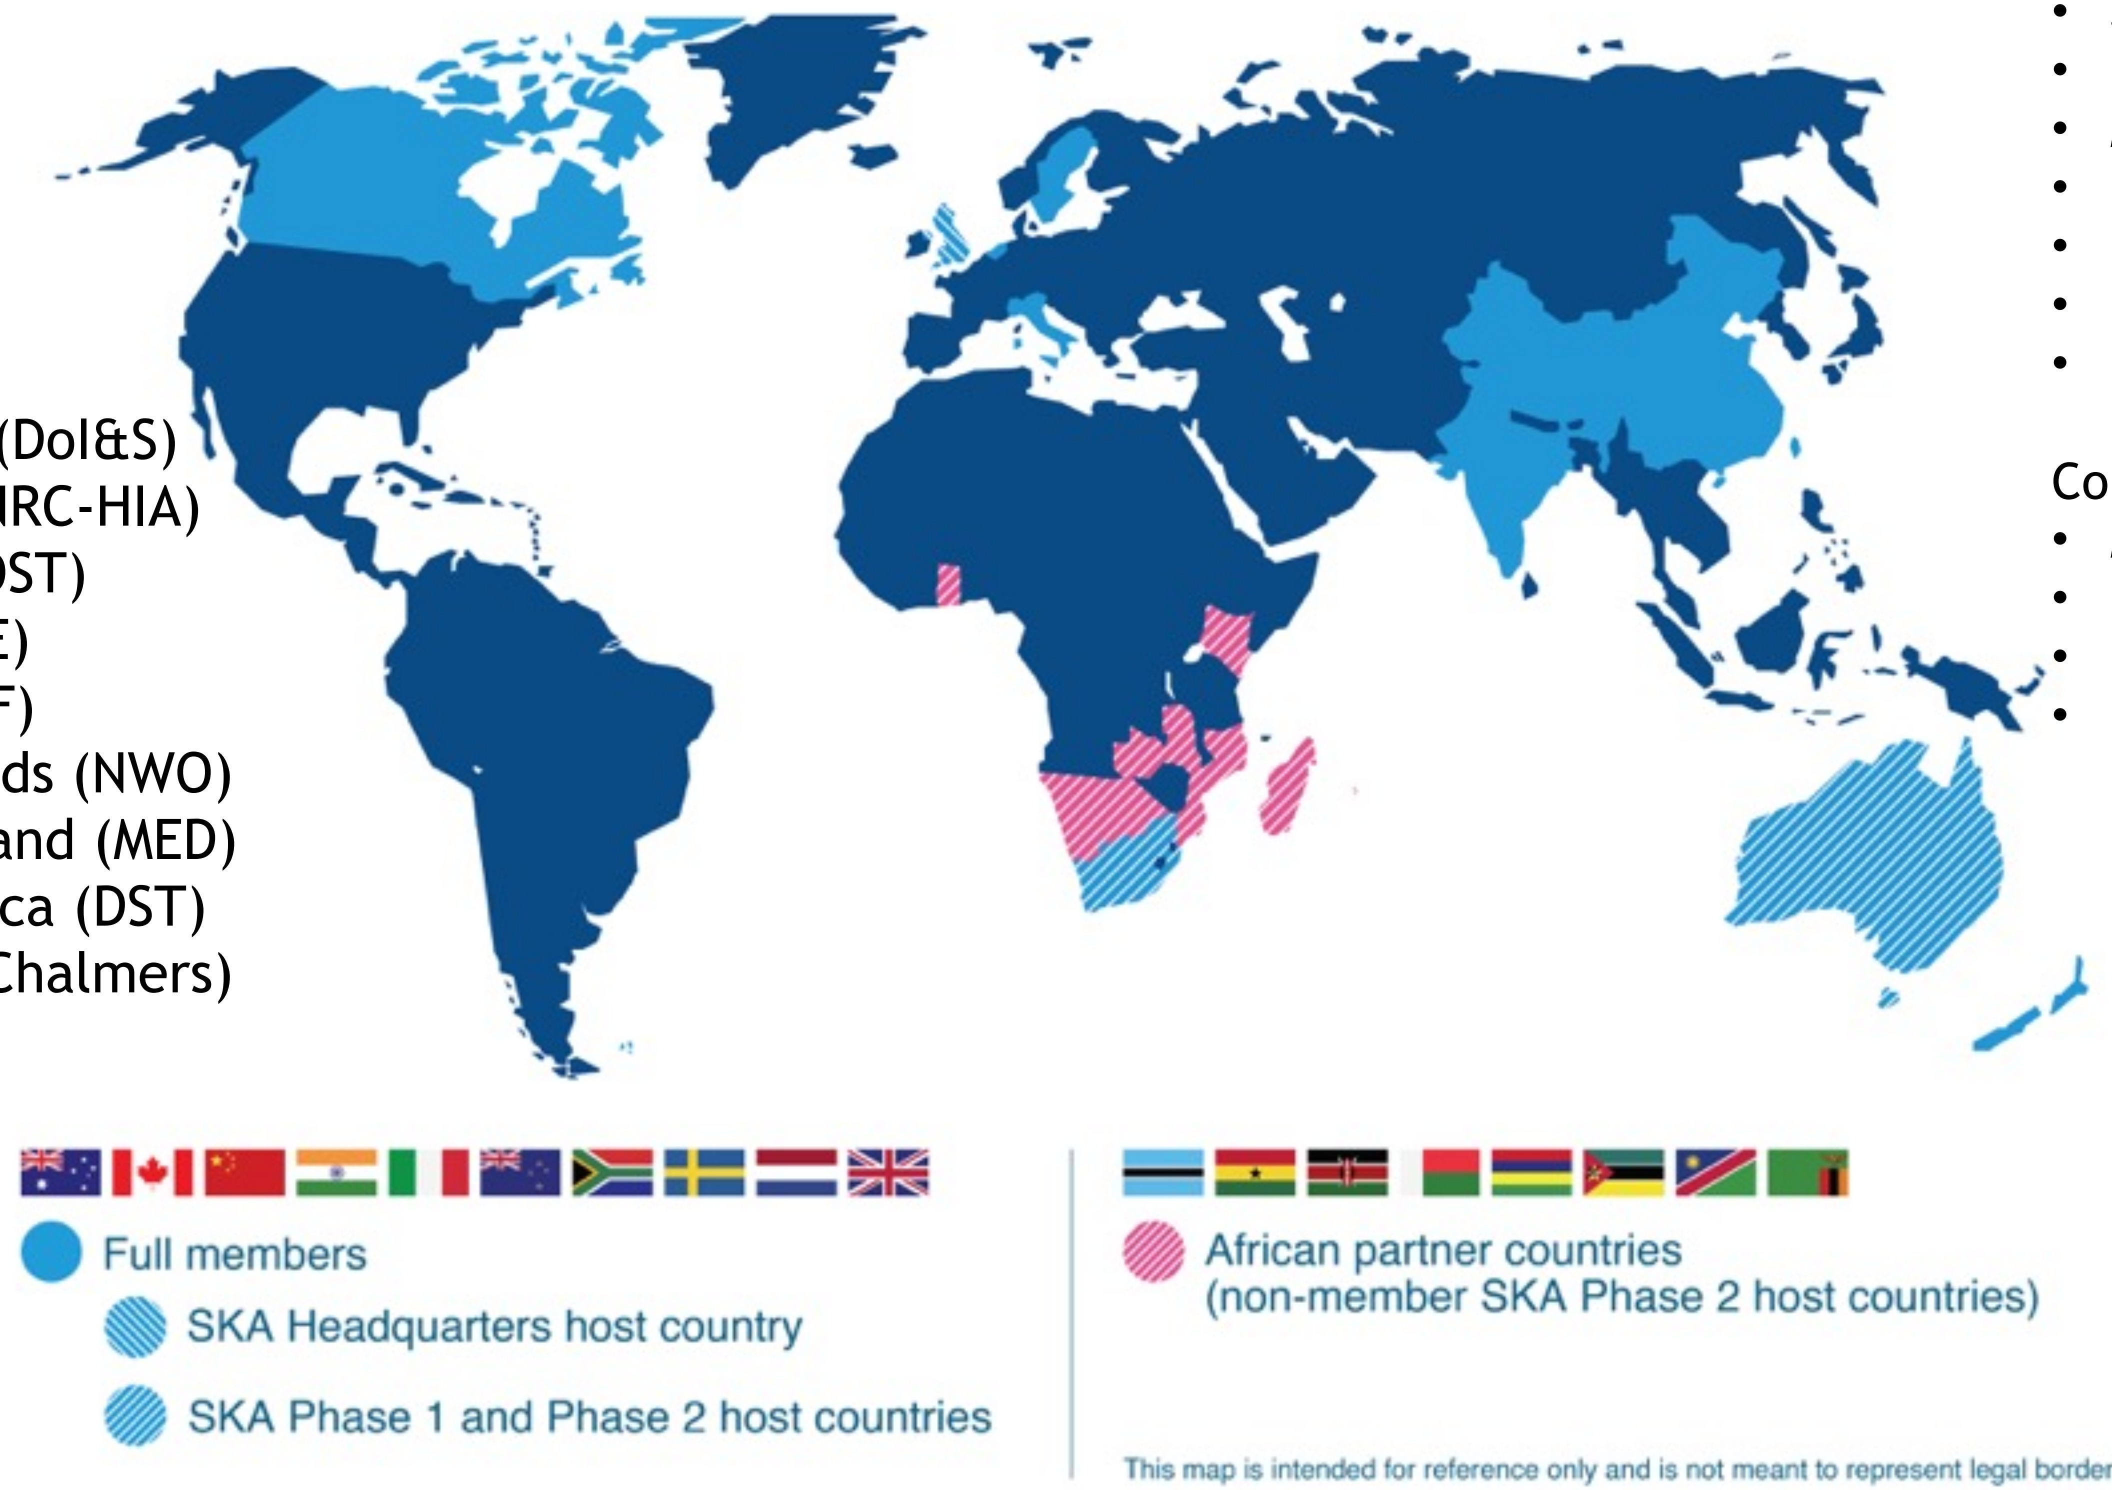

Interested Countries:

- France
- Germany
- Japan
- Korea
- Malta
- Portugal
- Spain
- Switzerland
- USA

Contacts:

- Mexico
- Brazil
- Ireland
- Russia

Phil Diamond Jan 2017

SKA Science Goals

Challenging Einstein

Galaxy Evolution, Cosmology and Dark Energy

Cosmic Magnetism

Cosmic Dawn

Cradle of Life

Continuum Surveys

Radio Transients

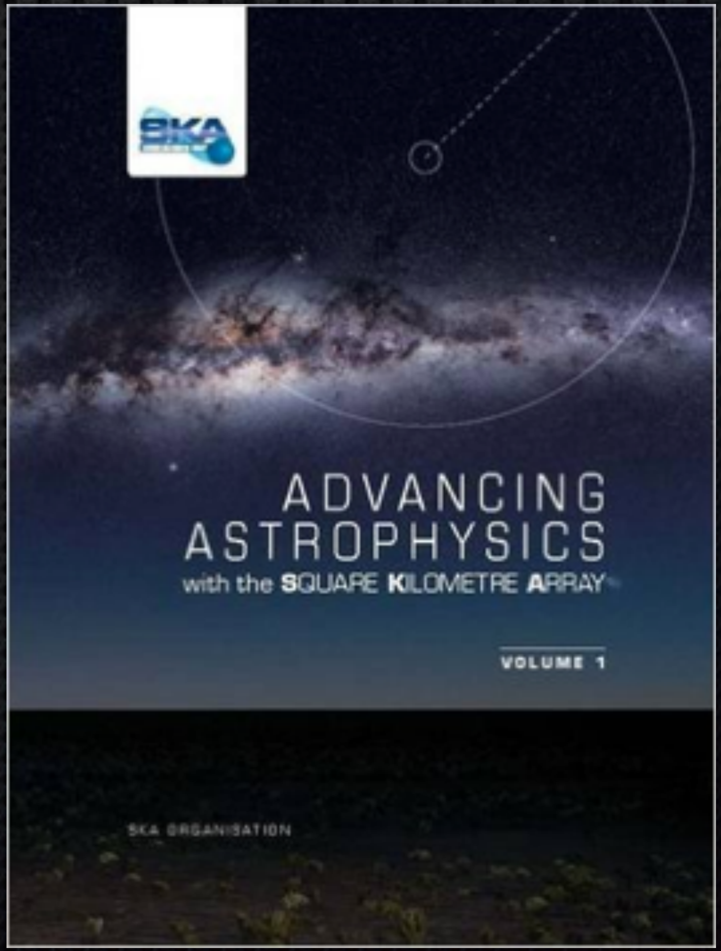

SKA1 Mid overview

- ~200, 13.5-m dishes
- 350 MHz — 14 GHz
- 2.5 GHz instantaneous bandwidth
- Compared to JVLA:
  - **4x** better resolution
  - **5x** more sensitivity
  - **60x** faster survey speed

SKA1 Low overview

- 131,072 dipole antennas (512 stations)
- 50 — 350 MHz
- 300 MHz instantaneous bandwidth
- Compared to LOFAR:
  - **25%** better resolution
  - **8x** more sensitivity
  - **125x** faster survey speed

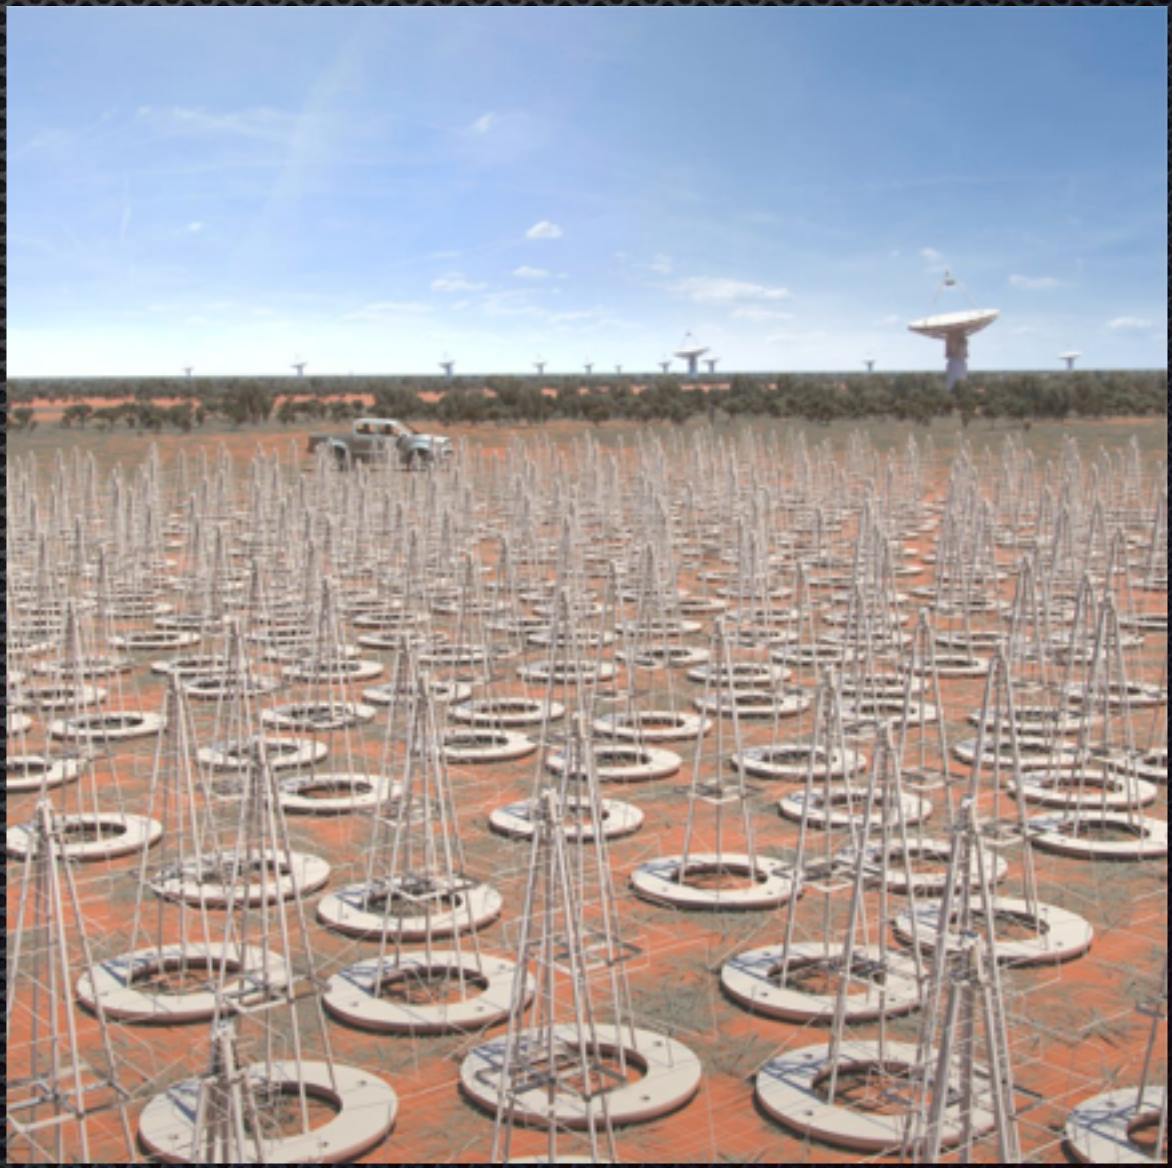

Big data challenges for the SKA

High-volume/velocity/dimensionality real-time processing

Telescope scheduling and process automation

Imaging and post-processing

Data storage and distribution

High-volume/velocity/dimensionality real-time processing

SKA1 Mid signal processing

digitise

80 Gb/s

80 Gb/s

80 Gb/s

16 Tb/s

SKA SQUARE KILOMETRE ARRAY CSP CENTRAL SIGNAL PROCESSOR

+197 more

High-volume/velocity/dimensionality real-time processing

SKA1 Mid signal processing

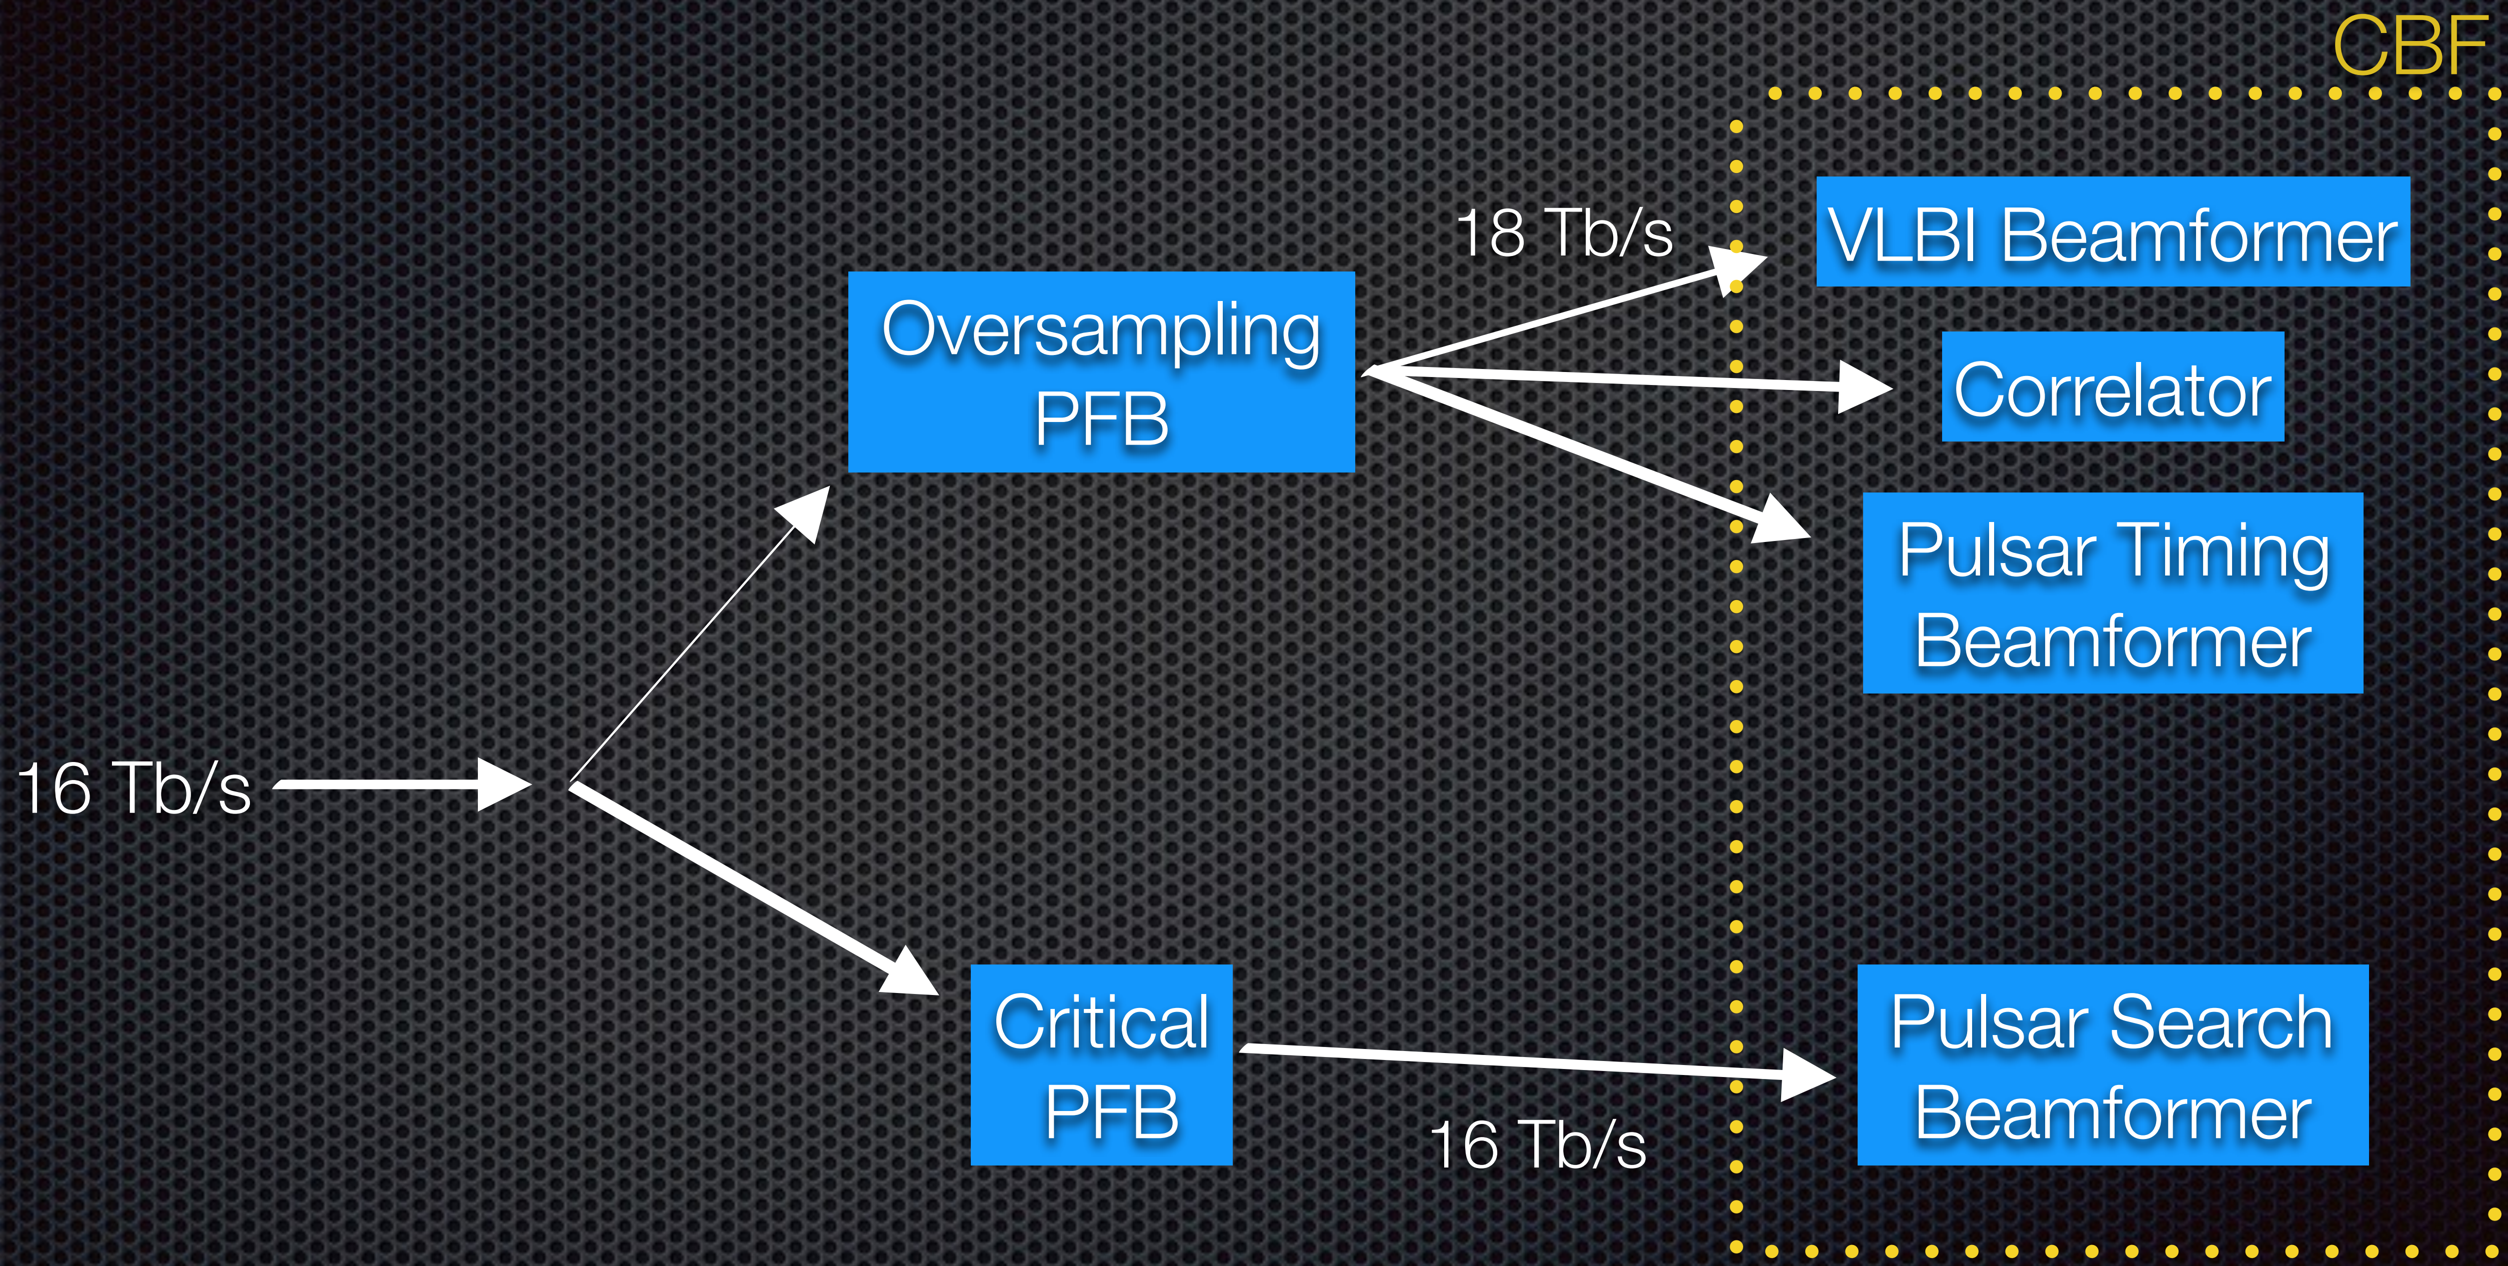

High-volume/velocity/dimensionality real-time processing

SKA1 Mid signal processing

CBF

VLBI Beamformer

Correlator

Pulsar Timing Beamformer

16 beams

Pulsar Search Beamformer

1500 beams

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

Pulsar Timing Engine

Pulsar Search Engine

Requires 10 PFlops/s

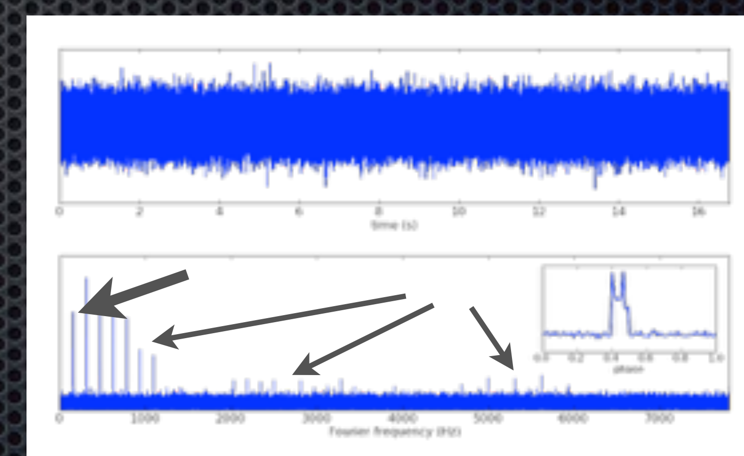

High-volume/velocity/dimensionality real-time processing

SKA1 Mid signal processing

Where's the challenge?

- Located in Karoo desert — logistically difficult
- Power cap ~700 kW
- Requires upwards of 15 PFlops/s
- Solution - Hybrid FGPA / GPU compute units

Telescope automation

- Each telescope supports up to 16 subarrays
- Multiple projects run simultaneously
- Each project provides constraints:
  - Sensitivity
  - Frequency band
  - Source position and epoch
- Further constraints imposed by:
  - Weather
  - Transient interferences sources (satellites, aircraft, etc.)

Combinatorial optimisation problem:

Constraint programming?

Ant Colony optimisation?

?

Imaging & post-processing

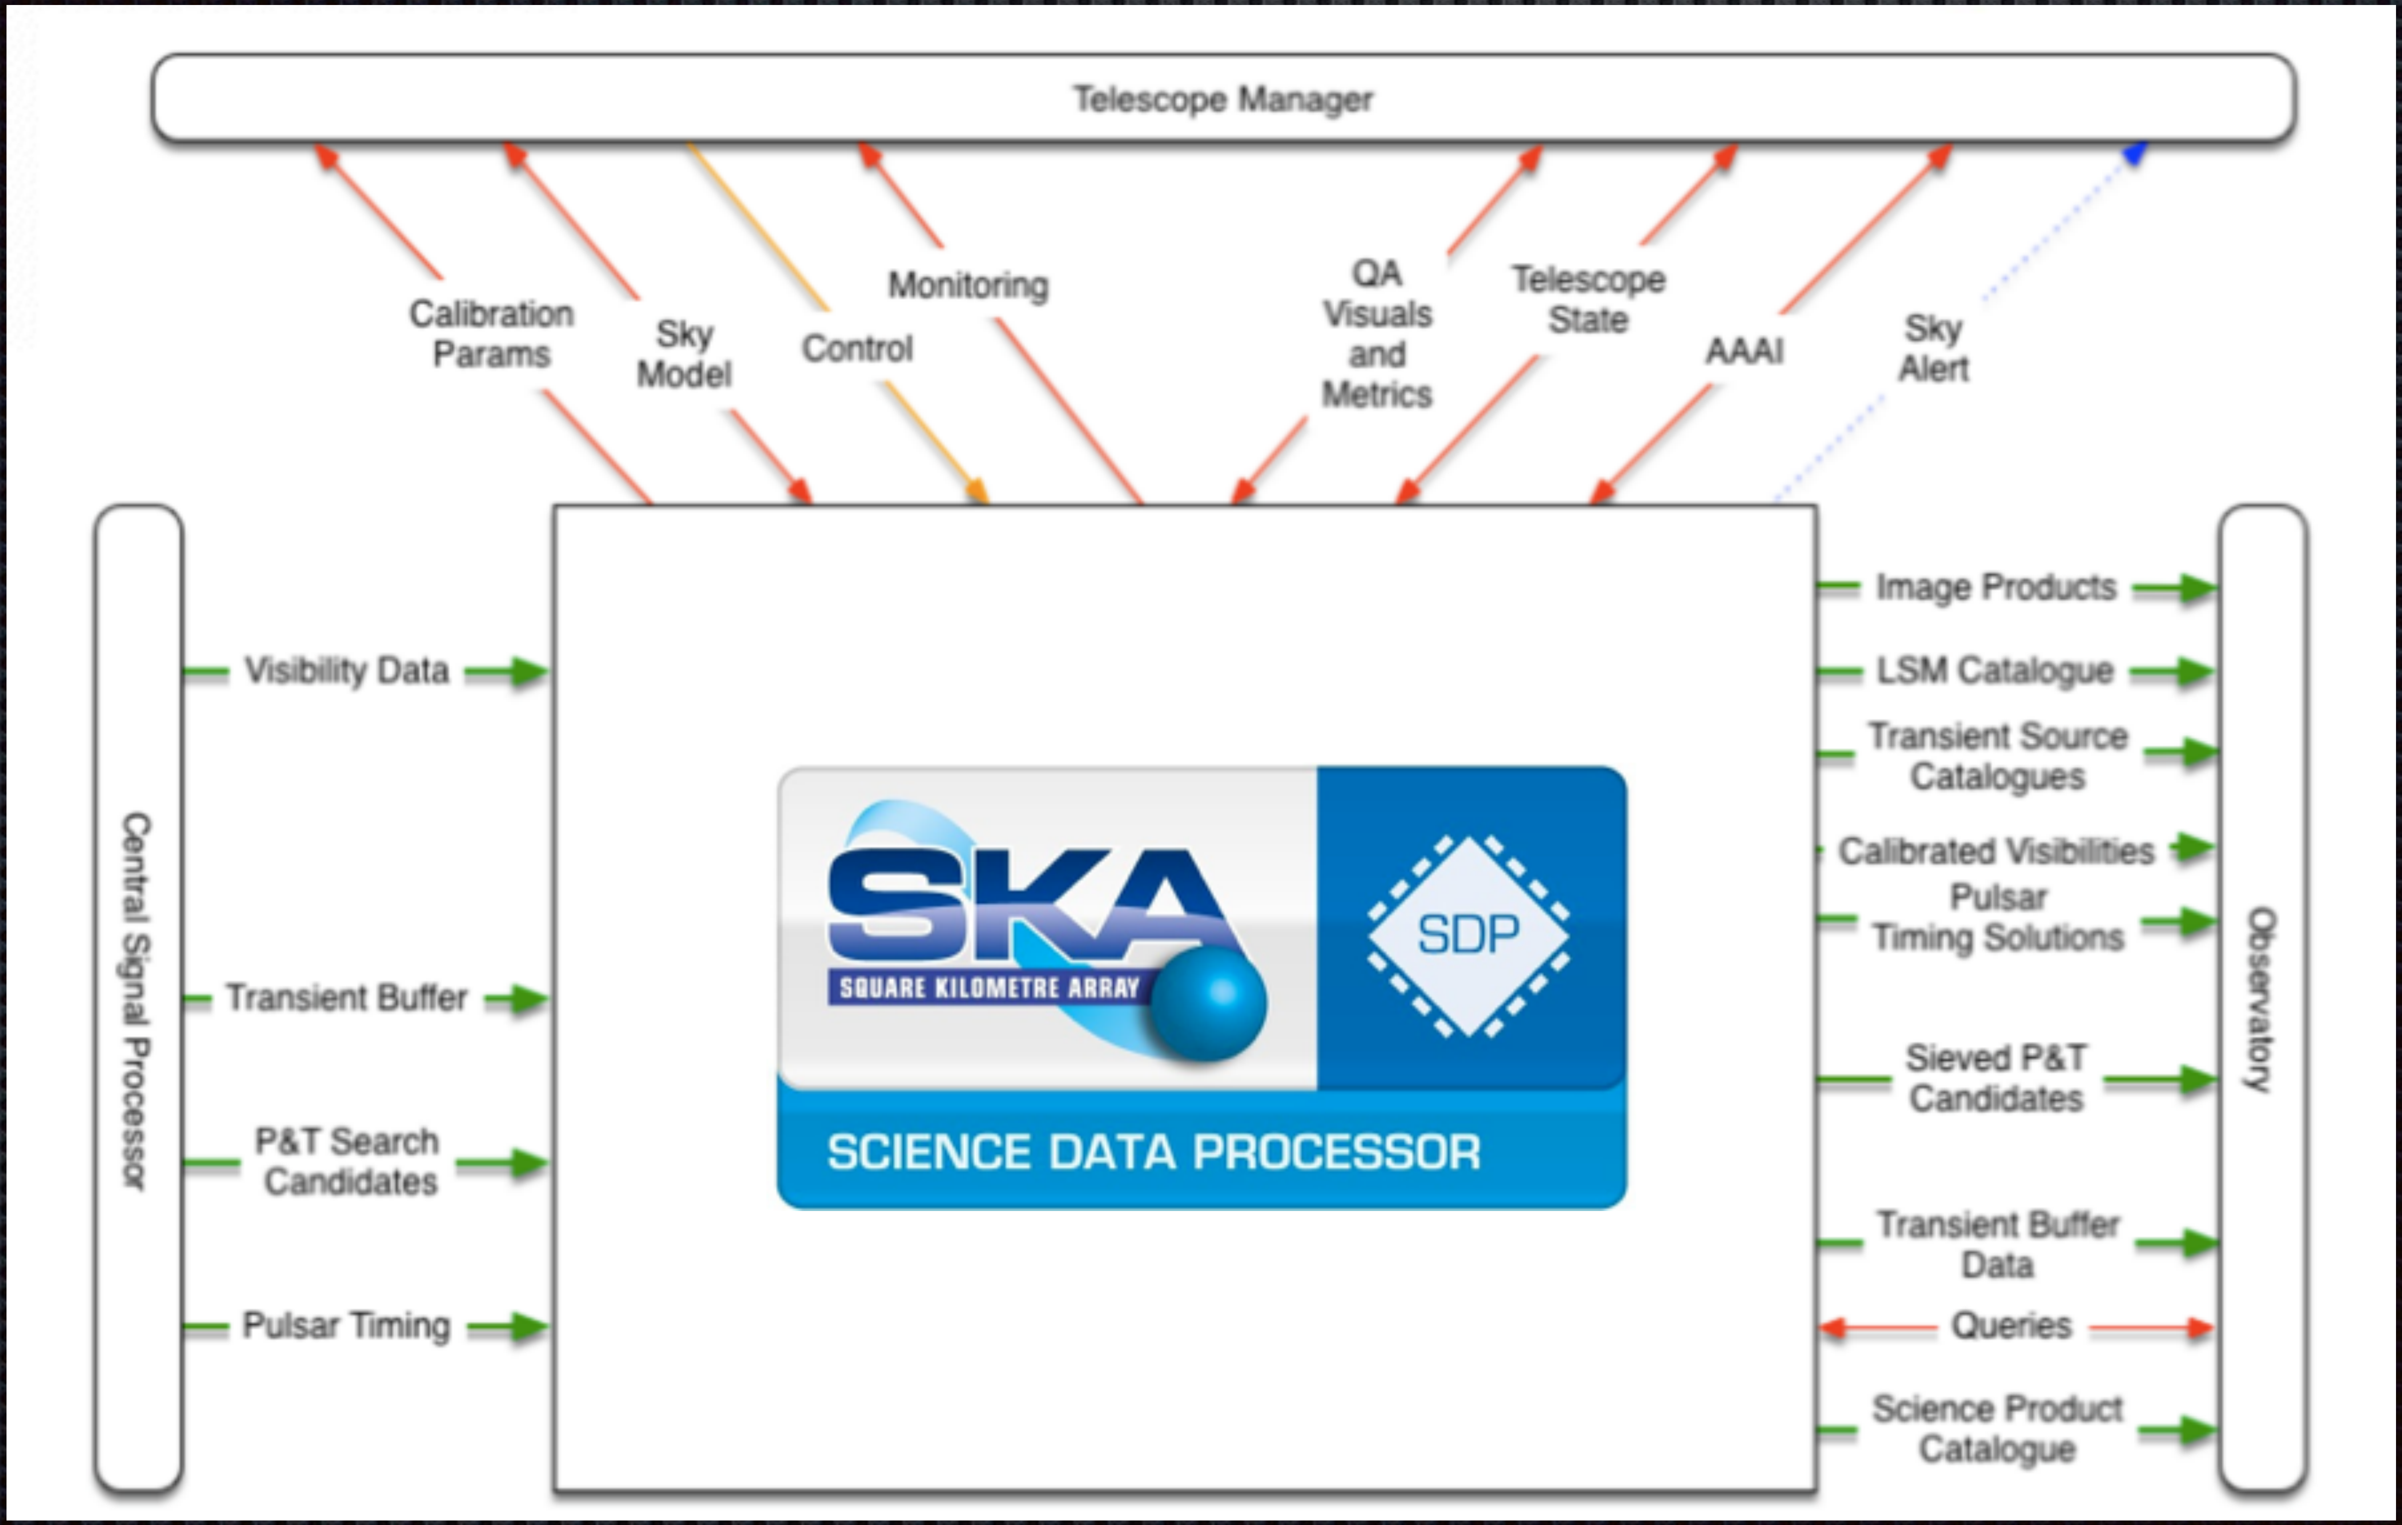

Cred. SDP, Anna Scaife

Imaging & post-processing

- SDP performs wide variety of tasks:
  - Imaging
  - Real-time calibration
  - Pulsar search candidate classification
  - Pulsar timing data reduction
- Where's the challenge?
  - Imaging requires **200-300 PFlops/s**!!!
  - 100 PB deep buffers (sustained 30 TB/s IO rate)
  - Power cap of 14 MW

Most challenging

Least challenging

Pulsar candidate classification

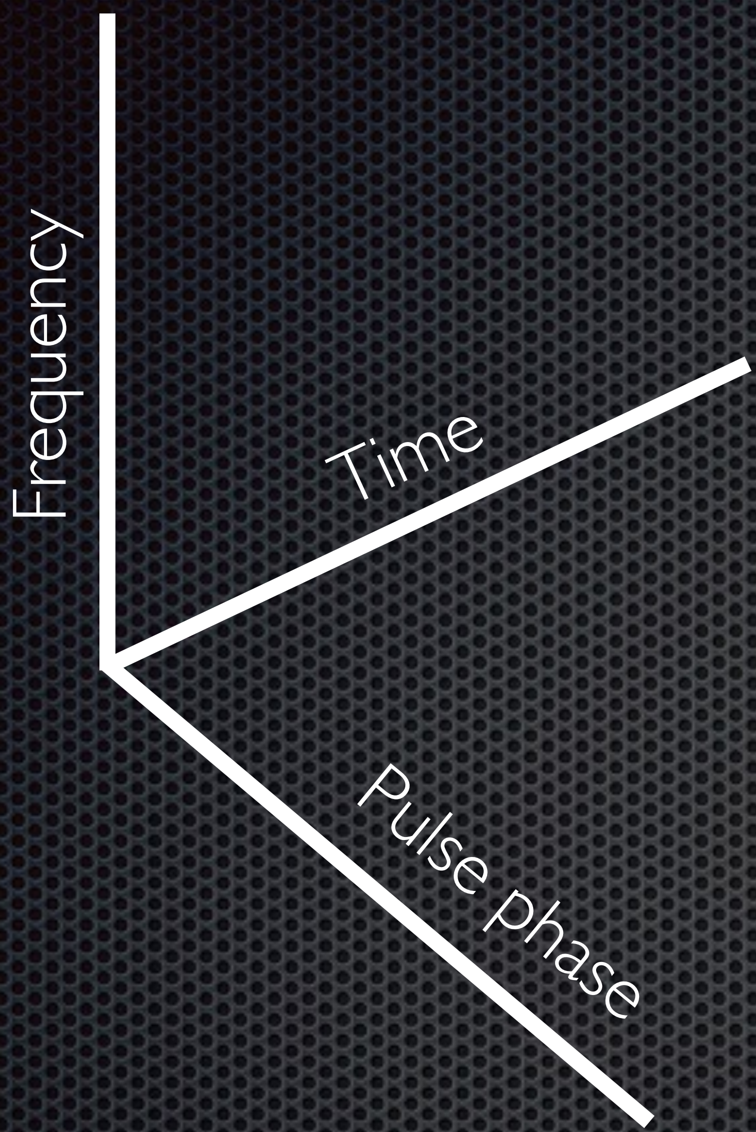

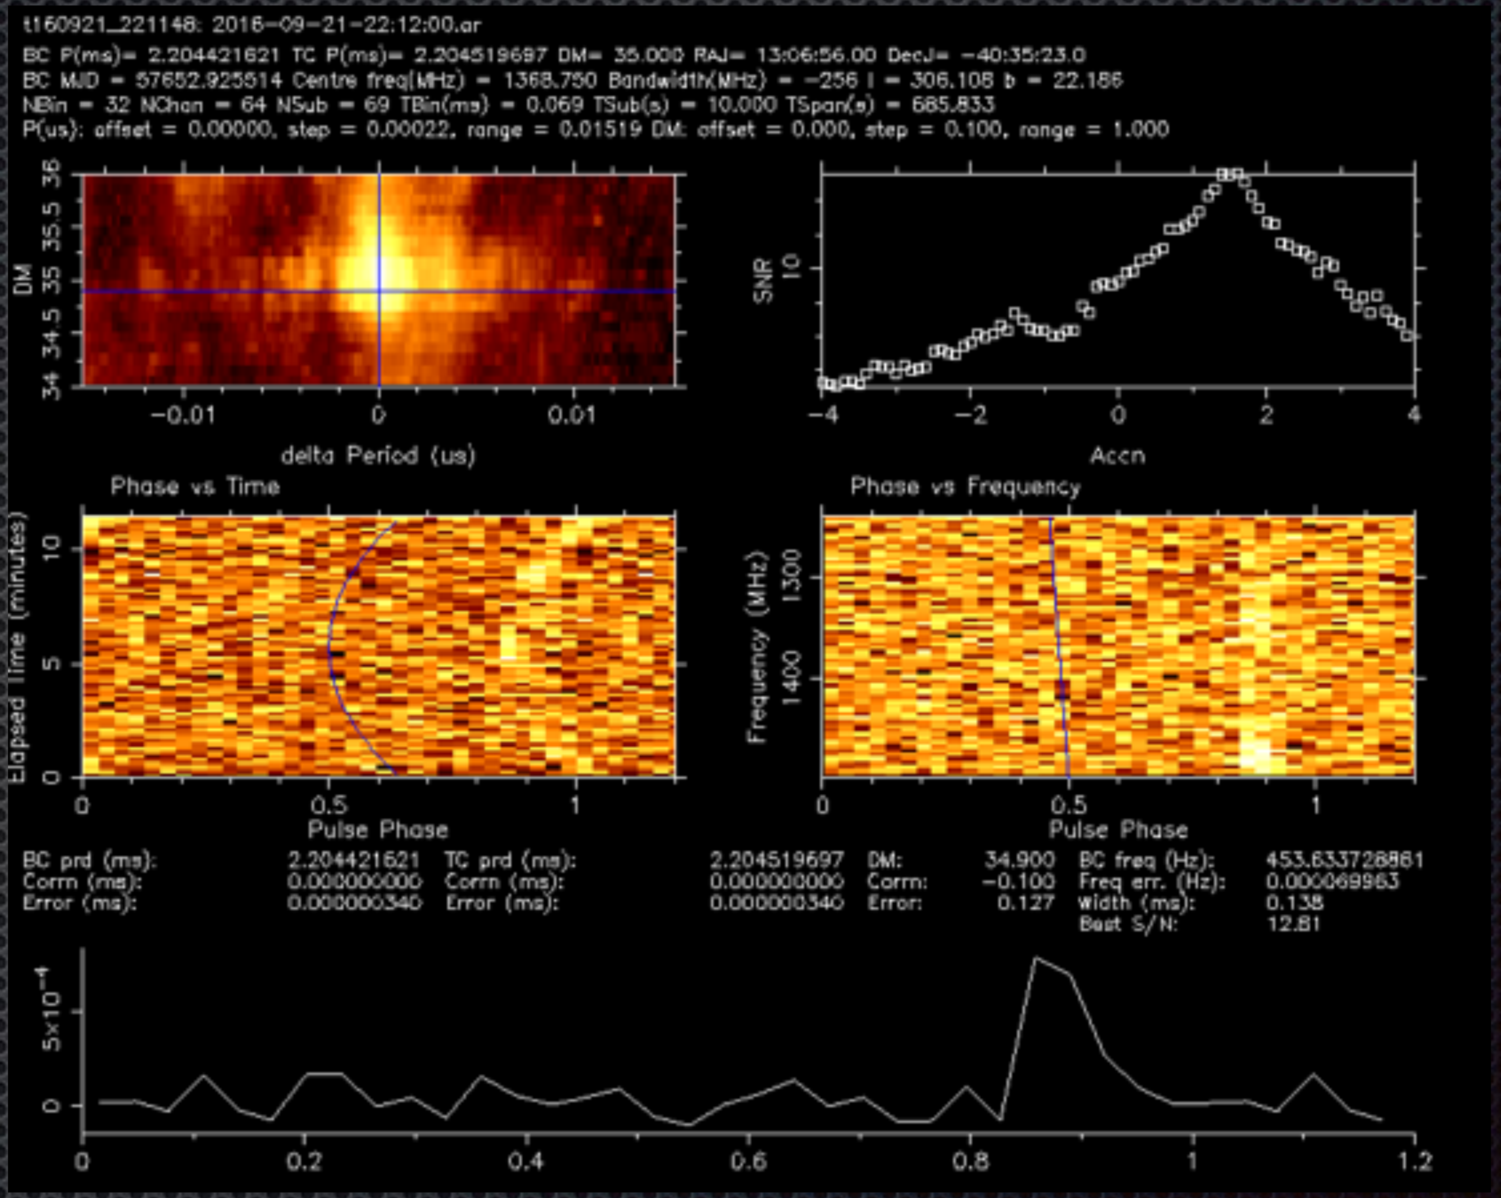

SKA1 Mid: 3000 candidates per second
56 billion per survey
Expect 15,000 to be “real” signals

Pulsar candidate classification

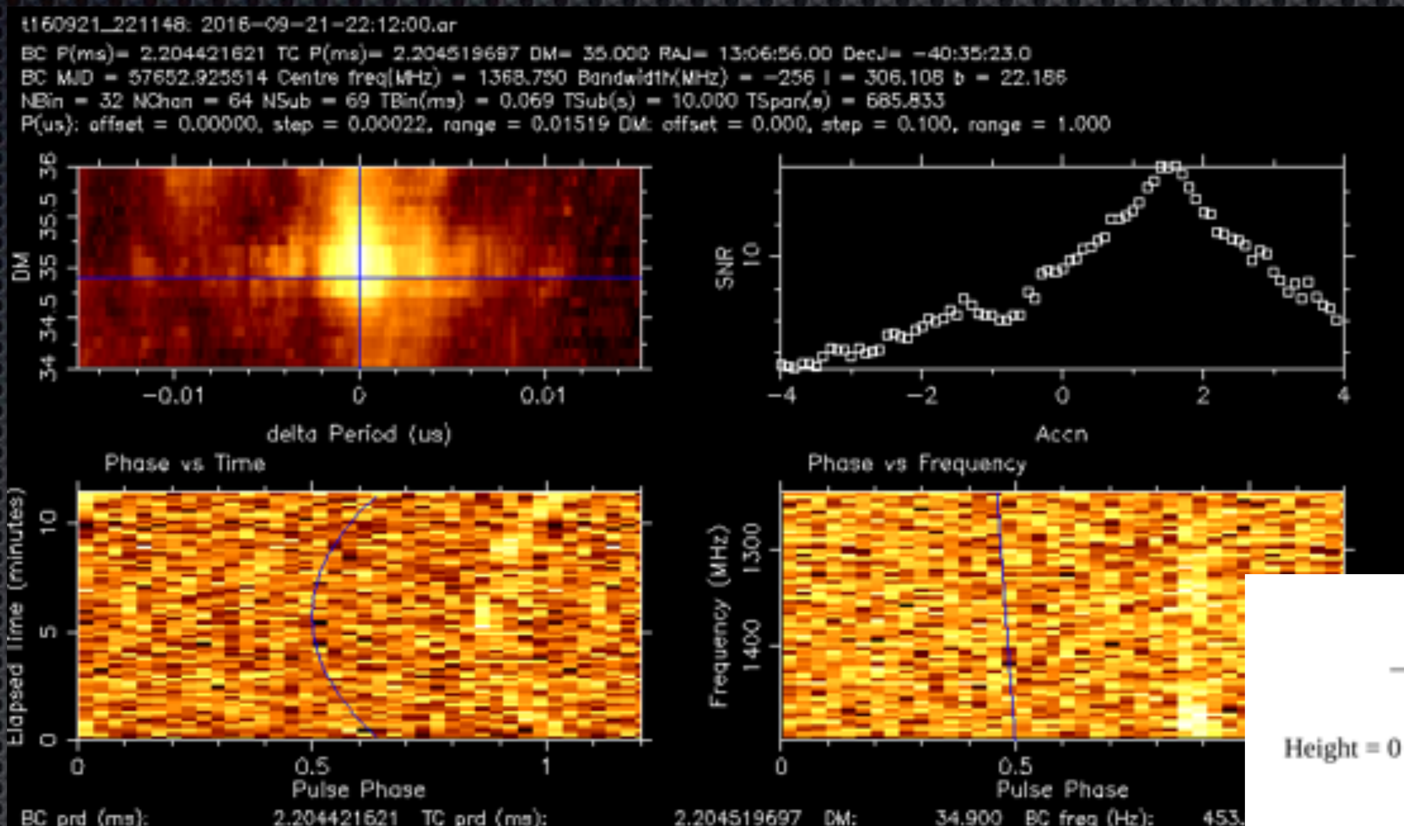

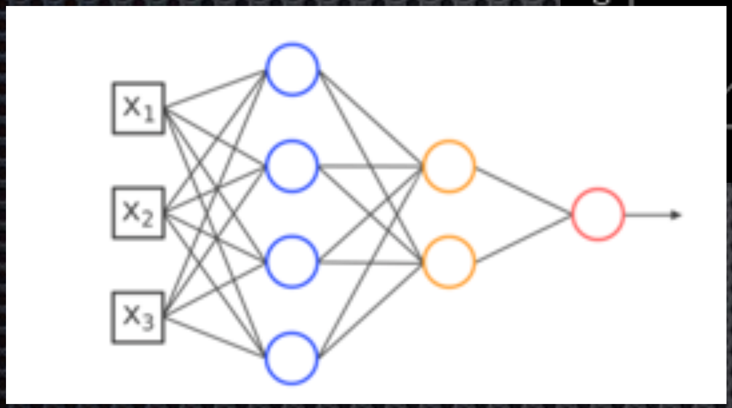

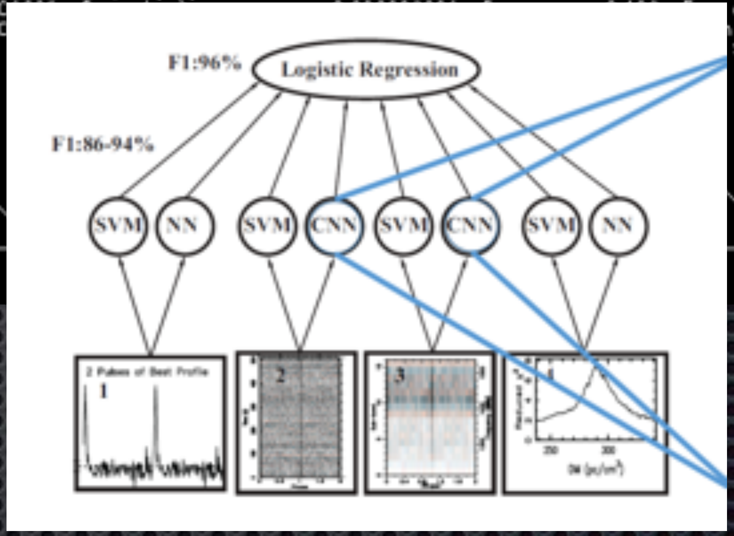

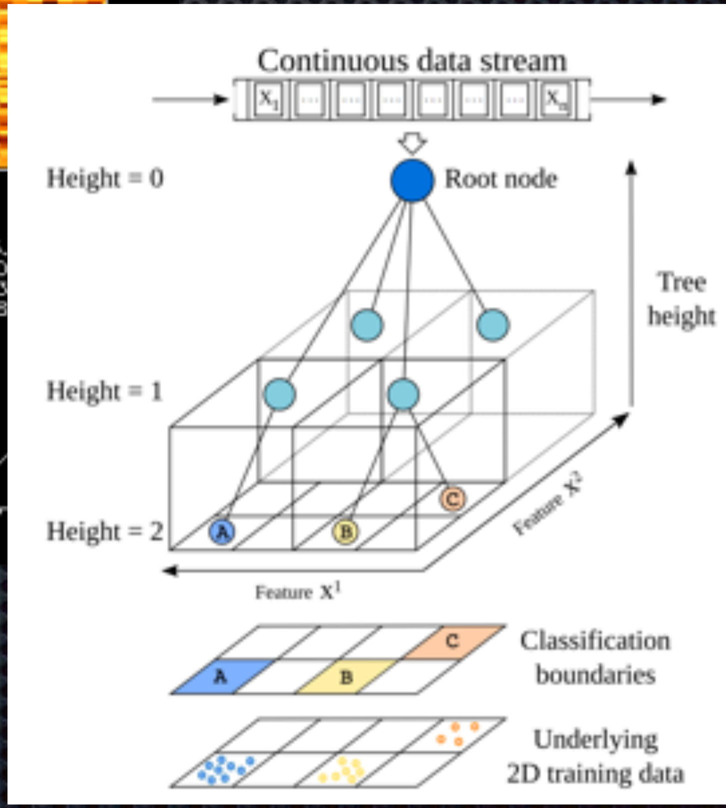

SPINN
Morello et al. (2014)

PICS
Zhu et al. (2013)

GH-FDT
Lyon et al. (2016)

Data storage & distribution

- 10-yr storage requirement for continuum survey of 5 EB
- High-speed international networks required to support data transfers to regional centres
- Regional centres provide support to local scientists
- This is (mostly) where the science will come from
- Data formats must be standardised (and self describing): FITS, HDF5, JSON, BSON, etc.

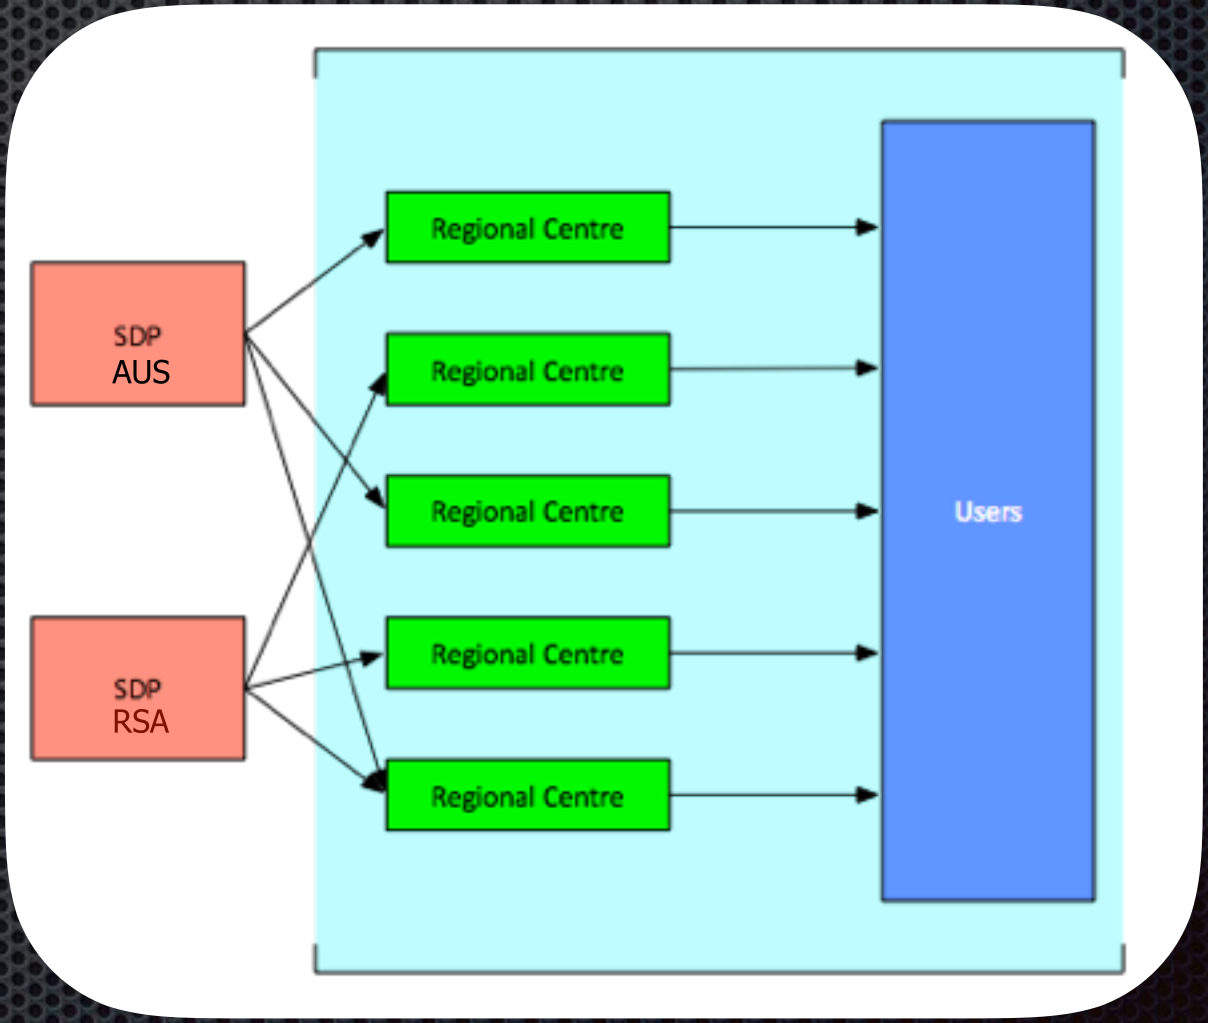

Custom Experiments

Precision measurements of cosmic ray air showers with the SKA

T. Huege[1], J.D. Bray[2], S. Buitink[3], R. Dallier[4,5], R.D. Ekers[6], H. Falcke[3,7], C.W. James[8], L. Martin[4,5], B. Revenu[4], O. Scholten[9] and F.G. Schröder[1]
[1]*KIT;* [2]*Univ. of Southampton;* [3]*Radboud Univ. Nijmegen;* [4]*Subatech, Nantes;* [5]*Station de radioastronomie de Nançay;* [6]*CSIRO ATNF;* [7]*ASTRON;* [8]*Univ. of Erlangen-Nuremberg;* [9]*Univ. of Groningen*
E-mail: tim.huege@kit.edu

Lunar detection of ultra-high-energy cosmic rays and neutrinos with the Square Kilometre Array

J.D. Bray*[1], J. Alvarez-Muñiz[2], S. Buitink[3], R.D. Dagkesamanskii[4], R.D. Ekers[5], H. Falcke[3,6], K.G. Gayley[7], T. Huege[8], C.W. James[9], M. Mevius[10], R.L. Mutel[7], R.J. Protheroe[11], O. Scholten[10], R.E. Spencer[12] and S. ter Veen[3]
[1]*Univ. of Southampton;* [2]*Univ. de Santiago de Compostela;* [3]*Radboud Univ. Nijmegen;* [4]*Lebedev Physical Institute;* [5]*CSIRO ATNF;* [6]*ASTRON;* [7]*Univ. of Iowa;* [8]*KIT;* [9]*Univ. of Erlangen-Nuremberg;* [10]*Univ. of Groningen;* [11]*Univ. of Adelaide;* [12]*Univ. of Manchester*
E-mail: j.bray@soton.ac.uk

- Particle detector co-location
- High-dynamic range deep transient buffers on each dipole (or subset of dipoles)
- Near-field imaging and beamforming
- Order arcsecond resolution

- Real-time search for nano-second scale radio bursts
- Similar to standard transient detection pipelines
- …“trivially” implemented in existing SKA hardware
- Synthesis filterbanks allow us to go back to Nyqvist sampling after beam forming

Custom Experiments

Pulsar Timing Engine

- Various custom experiment policies under discussion
- Simplest: Swappable processing pipelines on generic hardware!
- Docker is your friend…

- ECP-160056 - Custom Experiment Port - Type I - SDP interfaces
- ECP-160055 - Custom Experiment Port - Type I - CSP/PSS/PST interfaces
- ECP-160054 - Custom Experiment Port - Type II - CSP interfaces
- ECP-160053 - Custom Experiment Port - Type III - interfaces
- ECP-160052 - Custom Experiment Port - Type IV - interfaces to SKA-1 Low
- ECP-160051 - Custom Experiment Port - Type IV - interfaces to SKA-1 Mid
- ECP-160050 - Custom Experiment Port - Type V - interfaces to SKA-1 Low
- **ECP-160049 - Custom Experiment Port - Type V - interfaces to SKA-1 Mid**

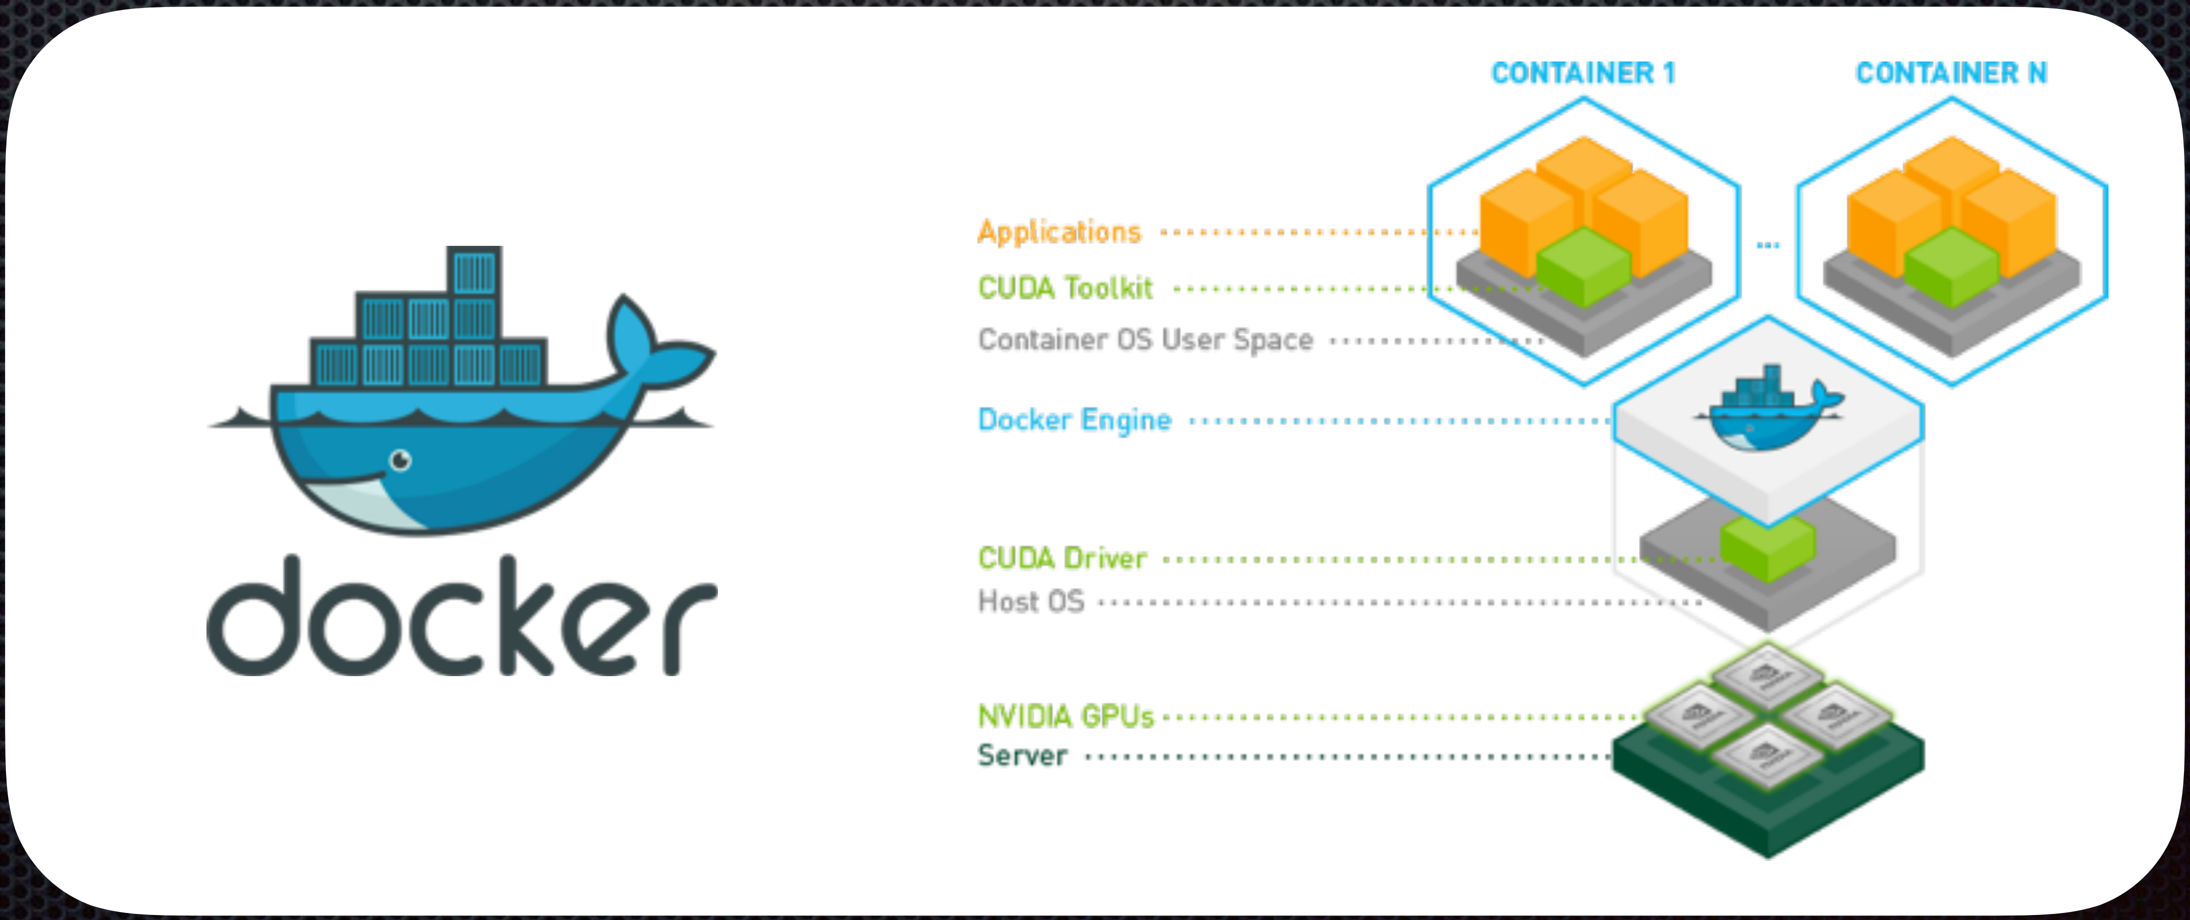

Precursor Instruments

MeerKAT

MWA

HERA

ASKAP

LOFAR

UTMOST

MeerKAT

Number of antennas	64 offset Gregorian
Dish diameter	13.5 m
Minimum baseline	29 m
Maximum baseline	8 km
Frequency bands (receivers)	0.58 – 1.015 GHz 0.9 – 1.67 GHz 8 – 14.5 GHz
Continuum imaging dynamic range at 1.4 GHz	50 dB
Line-to-line dynamic range at 1.4 GHz	43 dB
Mosaicing imaging dynamic range at 14 GHz	27 dB
Linear polarisation cross coupling across -3 dB beam	-30 dB
Sensitivity (0.58 – 1.67GHz)	220 m^2/K

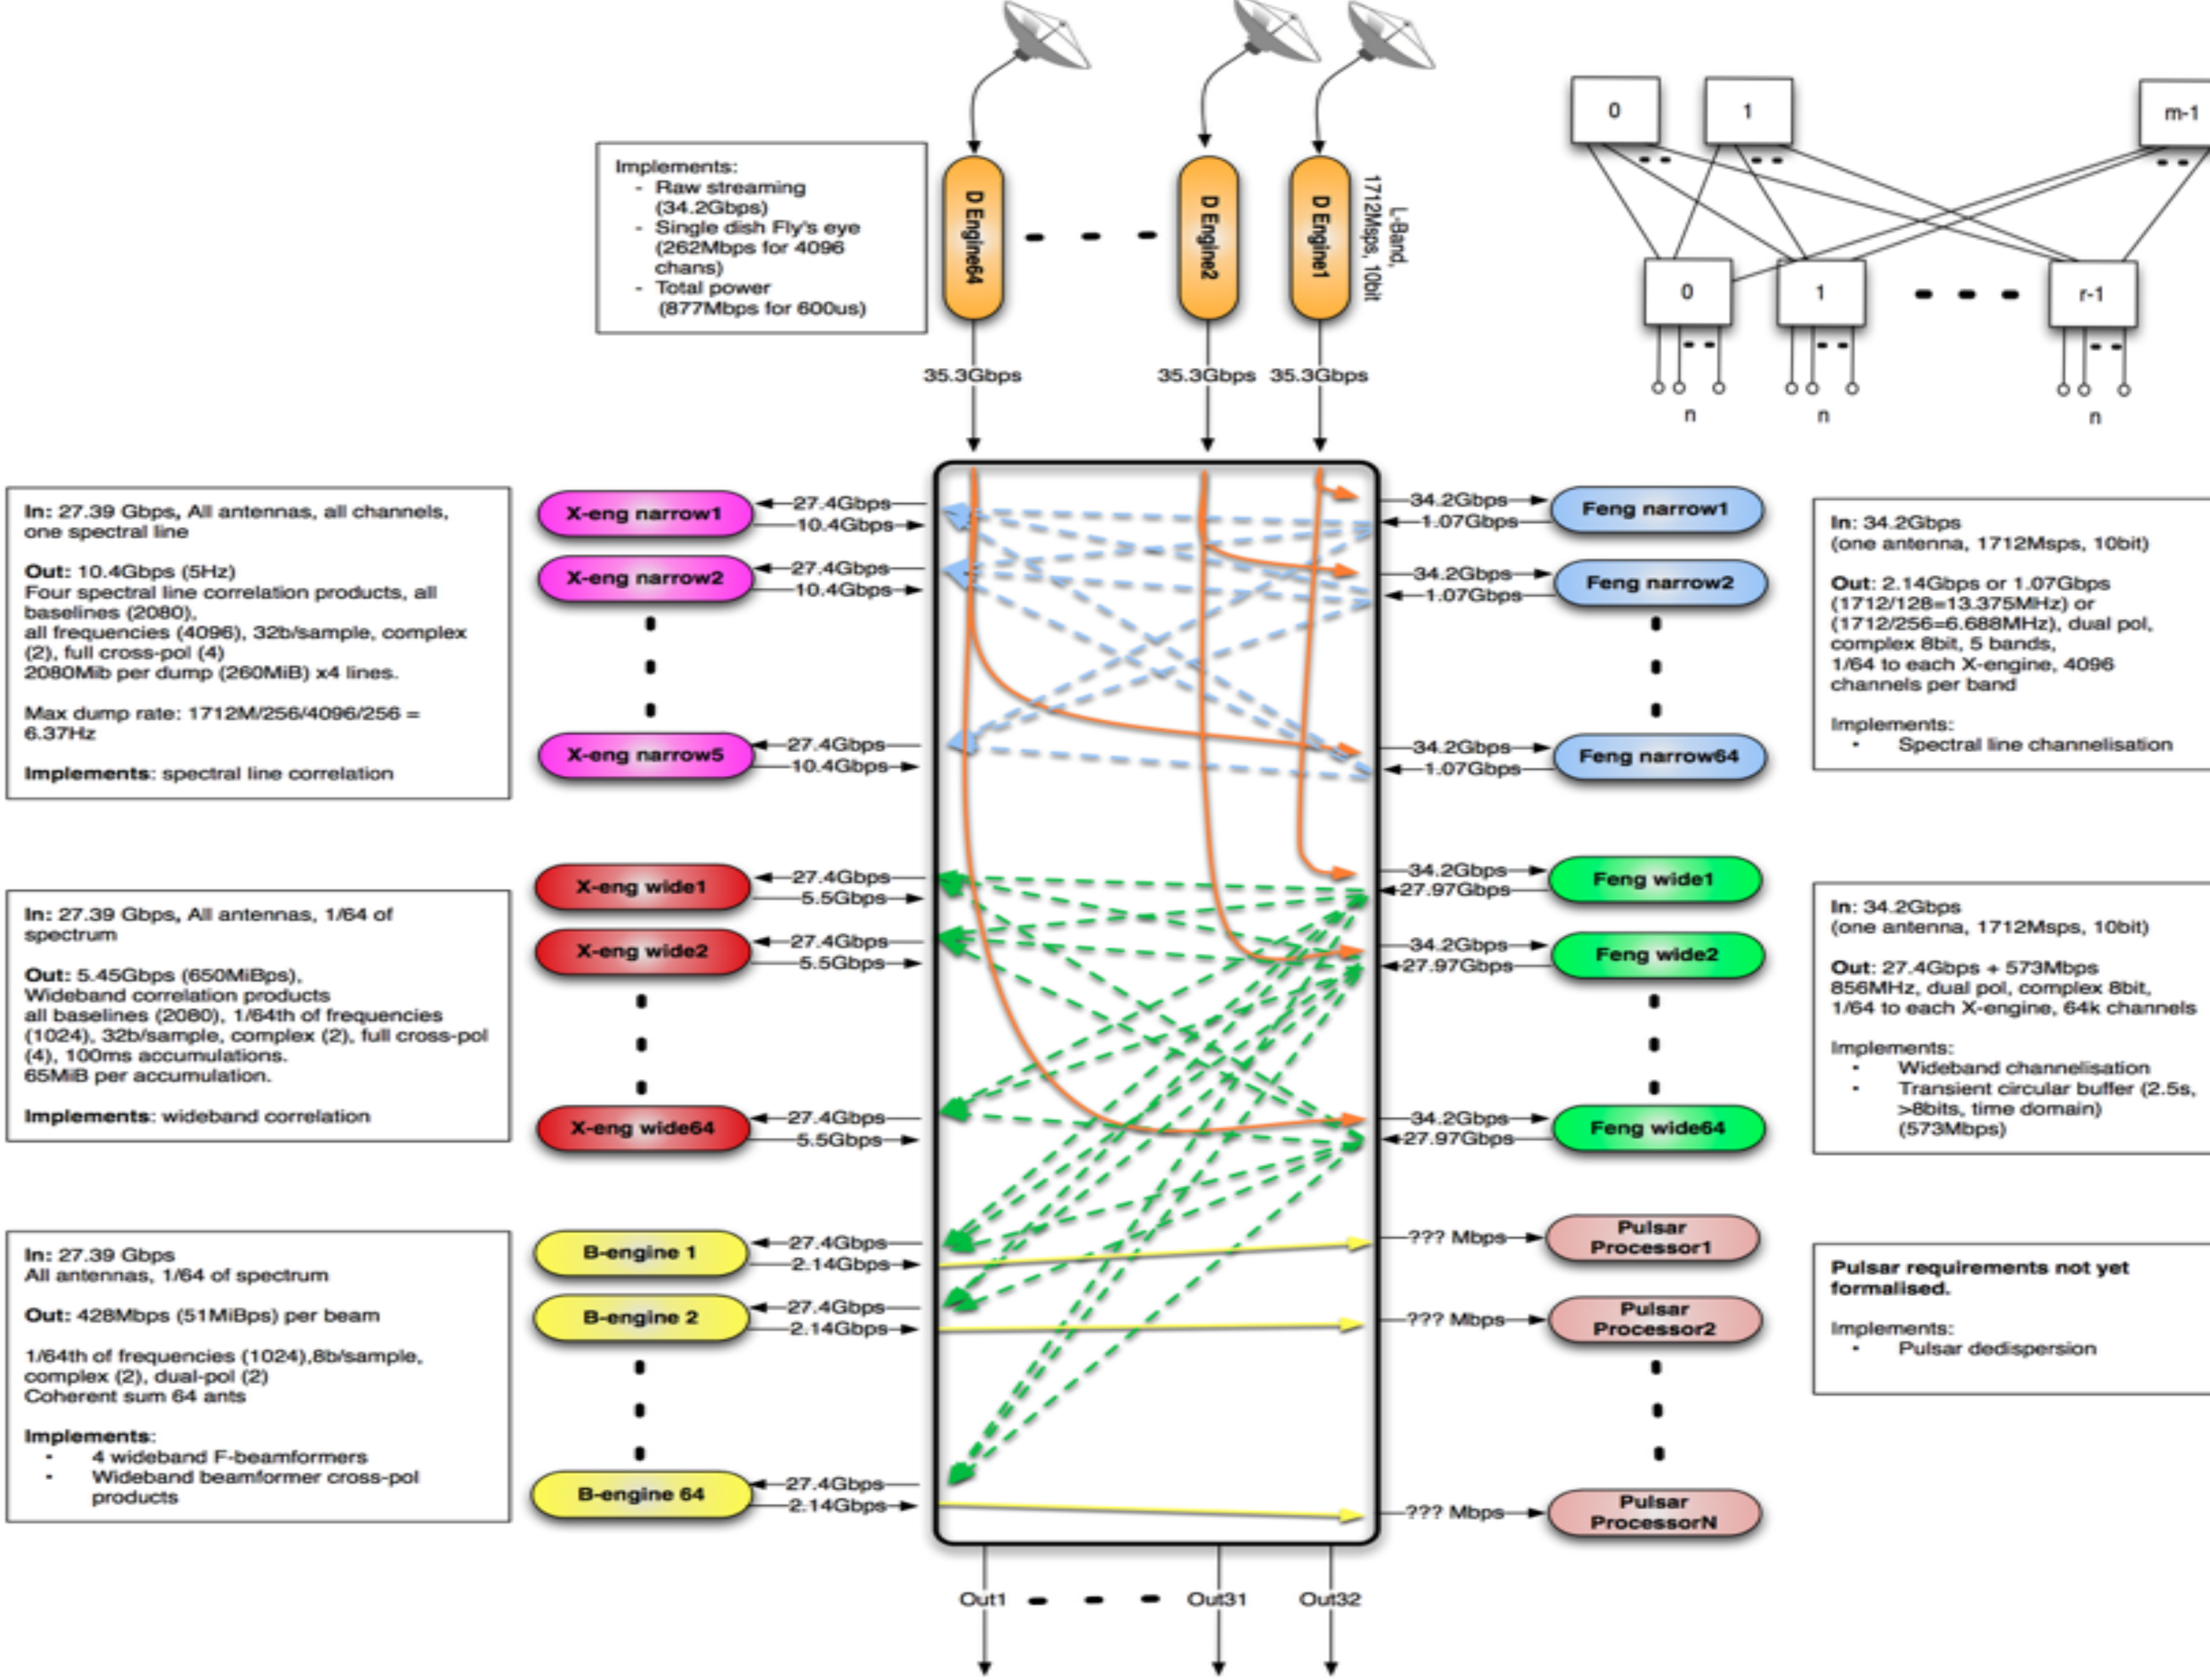

- Raw streaming (64x 34.2Gbps): 2191Gbps
- Correlator wide (100ms dumps): 358Gbps
- Temporally-averaged beamformer(100): 137Gbps
- Single dish channelised mode "Fly's eye": 56Gbps
- Beamformer: 27.4Gbps per beam (4): 55Gbps
- Correlator narrow (200ms dumps, 5 lines): 55Gbps
- VLBI beamformer (32 channels of 16MHz): 8Gbps
- Antennas voltage buffer: 5Gbps
- Incoherent sum: 1Gbps

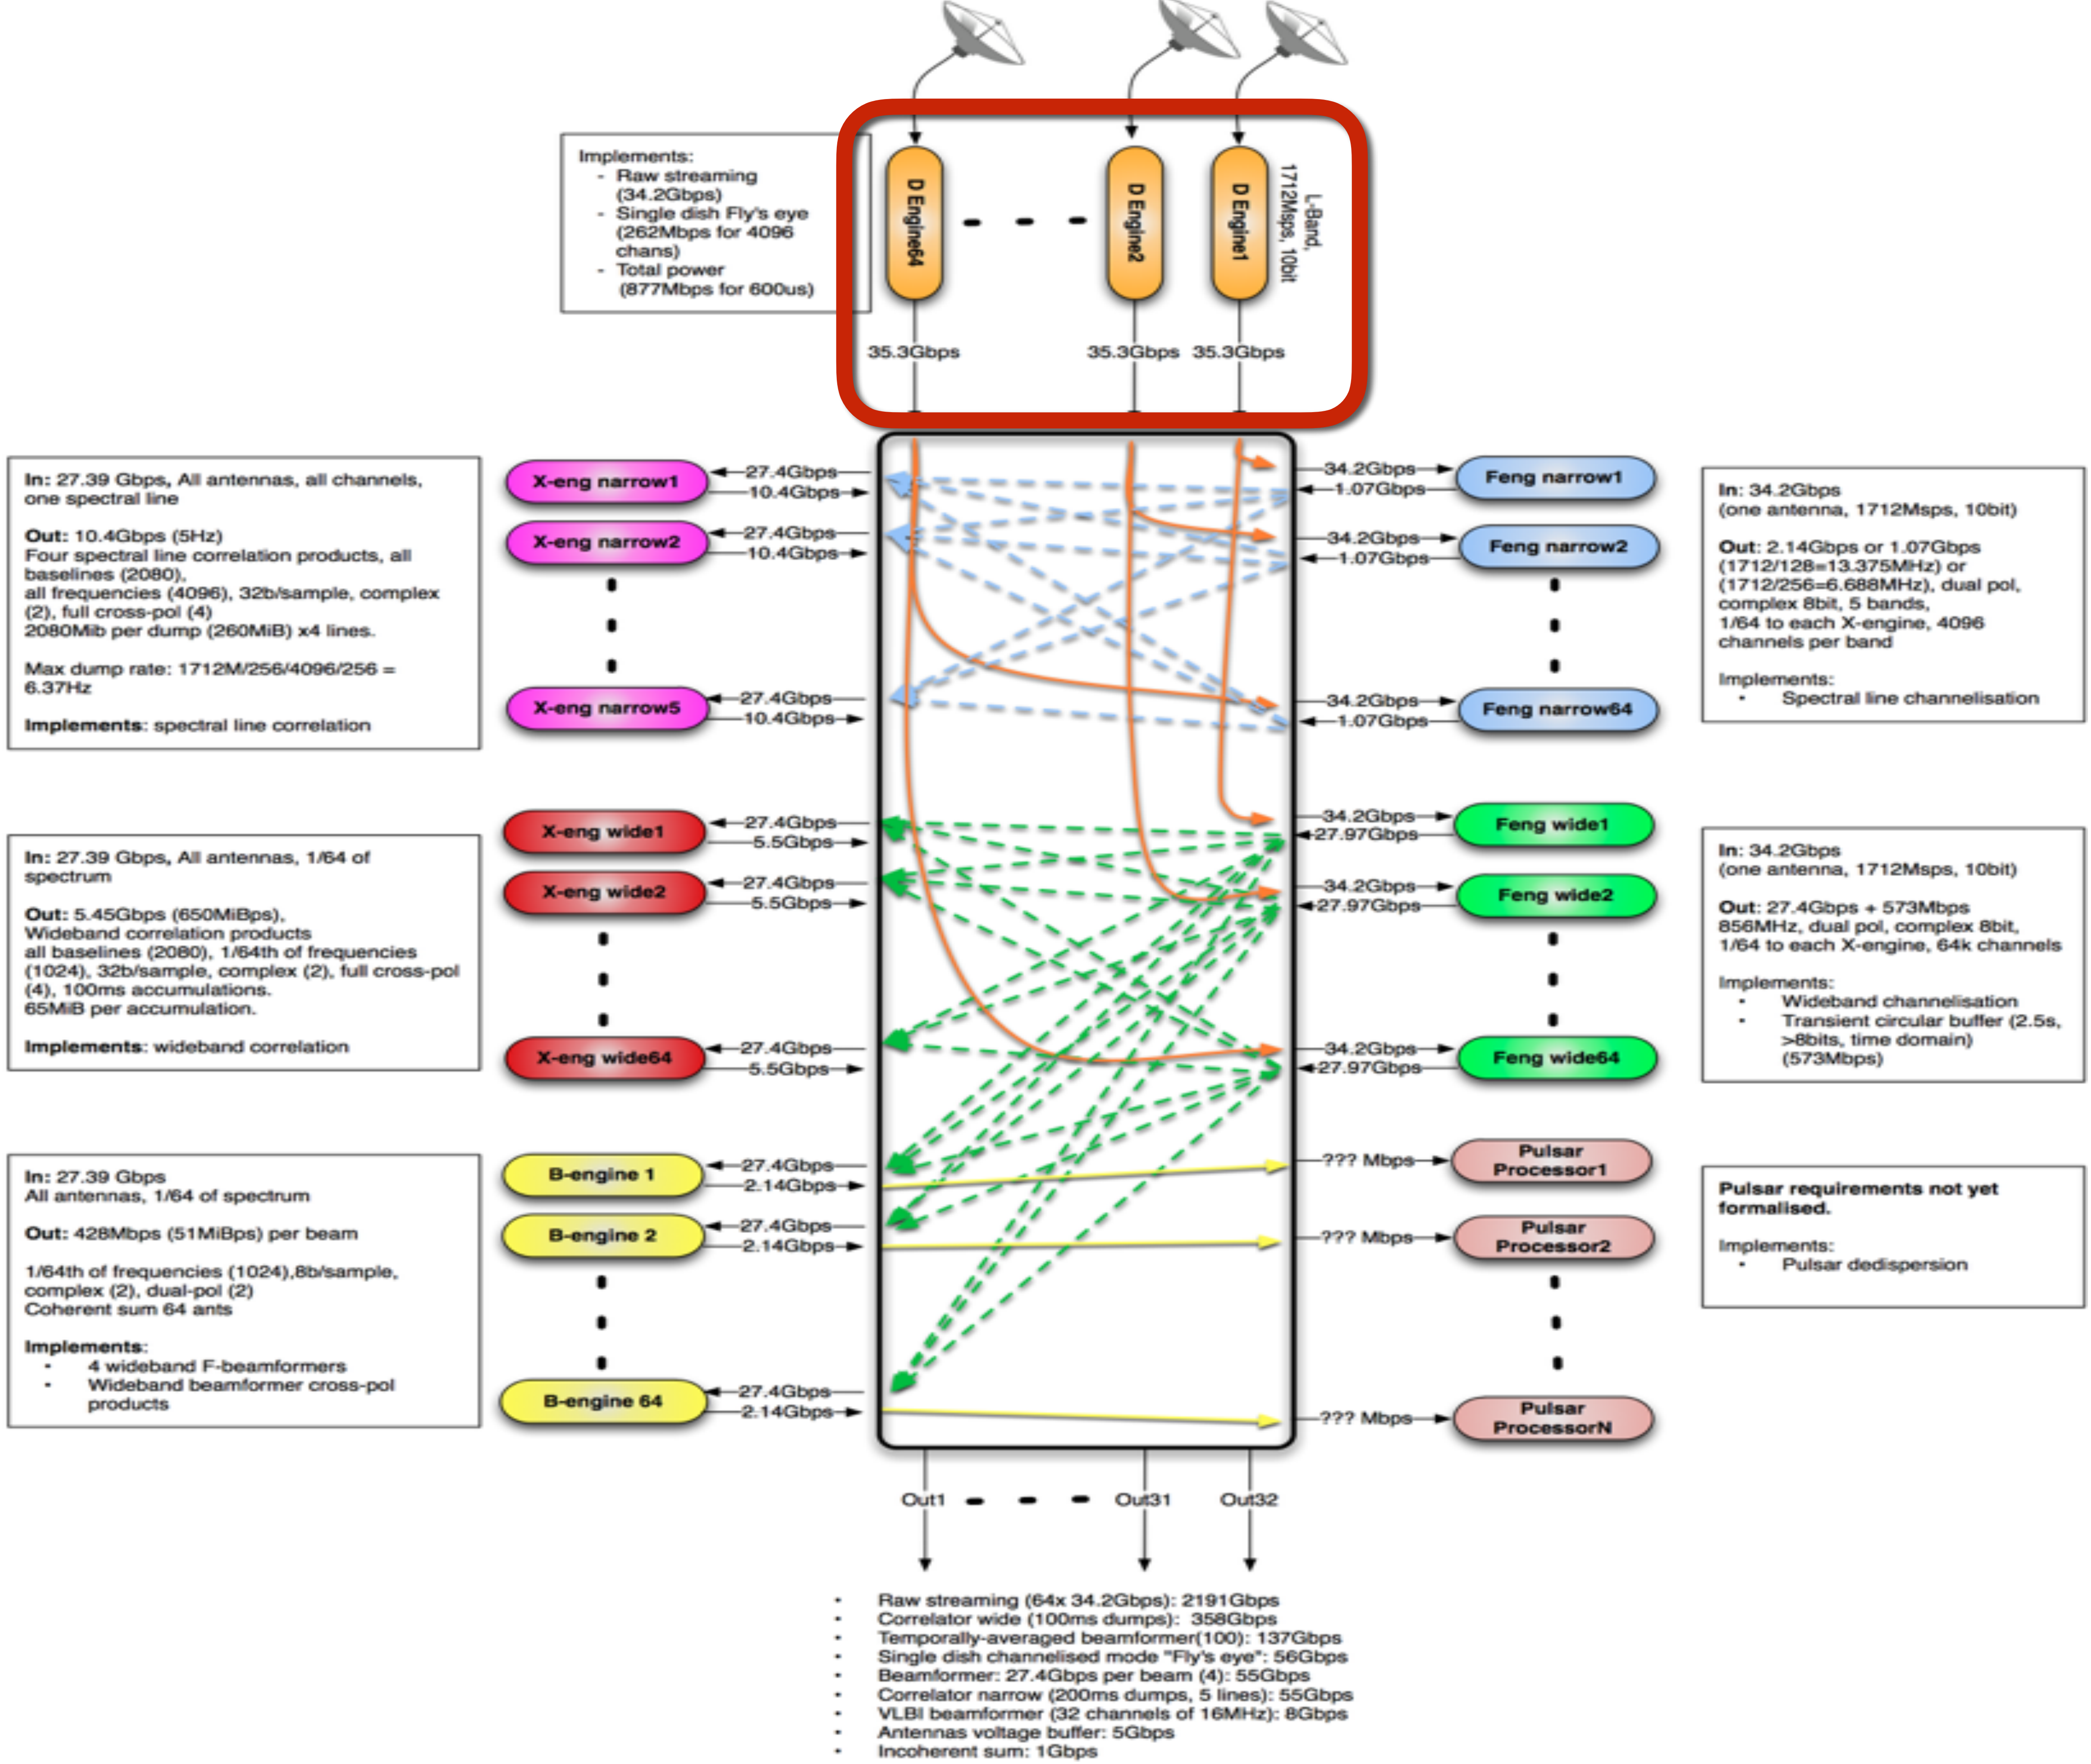
Implements:
- Raw streaming (34.2Gbps)
- Single dish Fly's eye (262Mbps for 4096 chans)
- Total power (877Mbps for 600us)
D Engine64
D Engine2
D Engine1
L-Band, 1712Msps, 10bit
35.3Gbps
35.3Gbps
35.3Gbps
In: 27.39 Gbps, All antennas, all channels, one spectral line
Out: 10.4Gbps (5Hz)
Four spectral line correlation products, all baselines (2080), all frequencies (4096), 32b/sample, complex (2), full cross-pol (4)
2080Mib per dump (260MiB) x4 lines.
Max dump rate: 1712M/256/4096/256 = 6.37Hz
Implements: spectral line correlation
X-eng narrow1
X-eng narrow2
X-eng narrow5
27.4Gbps
10.4Gbps
Feng narrow1
Feng narrow2
Feng narrow64
34.2Gbps
1.07Gbps
In: 34.2Gbps (one antenna, 1712Msps, 10bit)
Out: 2.14Gbps or 1.07Gbps (1712/128=13.375MHz) or (1712/256=6.688MHz), dual pol, complex 8bit, 5 bands, 1/64 to each X-engine, 4096 channels per band
Implements:
Spectral line channelisation
In: 27.39 Gbps, All antennas, 1/64 of spectrum
Out: 5.45Gbps (650MiBps), Wideband correlation products all baselines (2080), 1/64th of frequencies (1024), 32b/sample, complex (2), full cross-pol (4), 100ms accumulations.
65MiB per accumulation.
Implements: wideband correlation
X-eng wide1
X-eng wide2
X-eng wide64
27.4Gbps
5.5Gbps
Feng wide1
Feng wide2
Feng wide64
34.2Gbps
27.97Gbps
In: 34.2Gbps (one antenna, 1712Msps, 10bit)
Out: 27.4Gbps + 573Mbps 856MHz, dual pol, complex 8bit, 1/64 to each X-engine, 64k channels
Implements:
Wideband channelisation
Transient circular buffer (2.5s, >8bits, time domain) (573Mbps)
In: 27.39 Gbps
All antennas, 1/64 of spectrum
Out: 428Mbps (51MiBps) per beam
1/64th of frequencies (1024),8b/sample, complex (2), dual-pol (2)
Coherent sum 64 ants
Implements:
4 wideband F-beamformers
Wideband beamformer cross-pol products
B-engine 1
B-engine 2
B-engine 64
27.4Gbps
2.14Gbps
??? Mbps
Pulsar Processor1
Pulsar Processor2
Pulsar ProcessorN
Pulsar requirements not yet formalised.
Implements:
Pulsar dedispersion
Out1
Out31
Out32
Raw streaming (64x 34.2Gbps): 2191Gbps
Correlator wide (100ms dumps): 358Gbps
Temporally-averaged beamformer(100): 137Gbps
Single dish channelised mode "Fly's eye": 56Gbps
Beamformer: 27.4Gbps per beam (4): 55Gbps
Correlator narrow (200ms dumps, 5 lines): 55Gbps
VLBI beamformer (32 channels of 16MHz): 8Gbps
Antennas voltage buffer: 5Gbps
Incoherent sum: 1Gbps

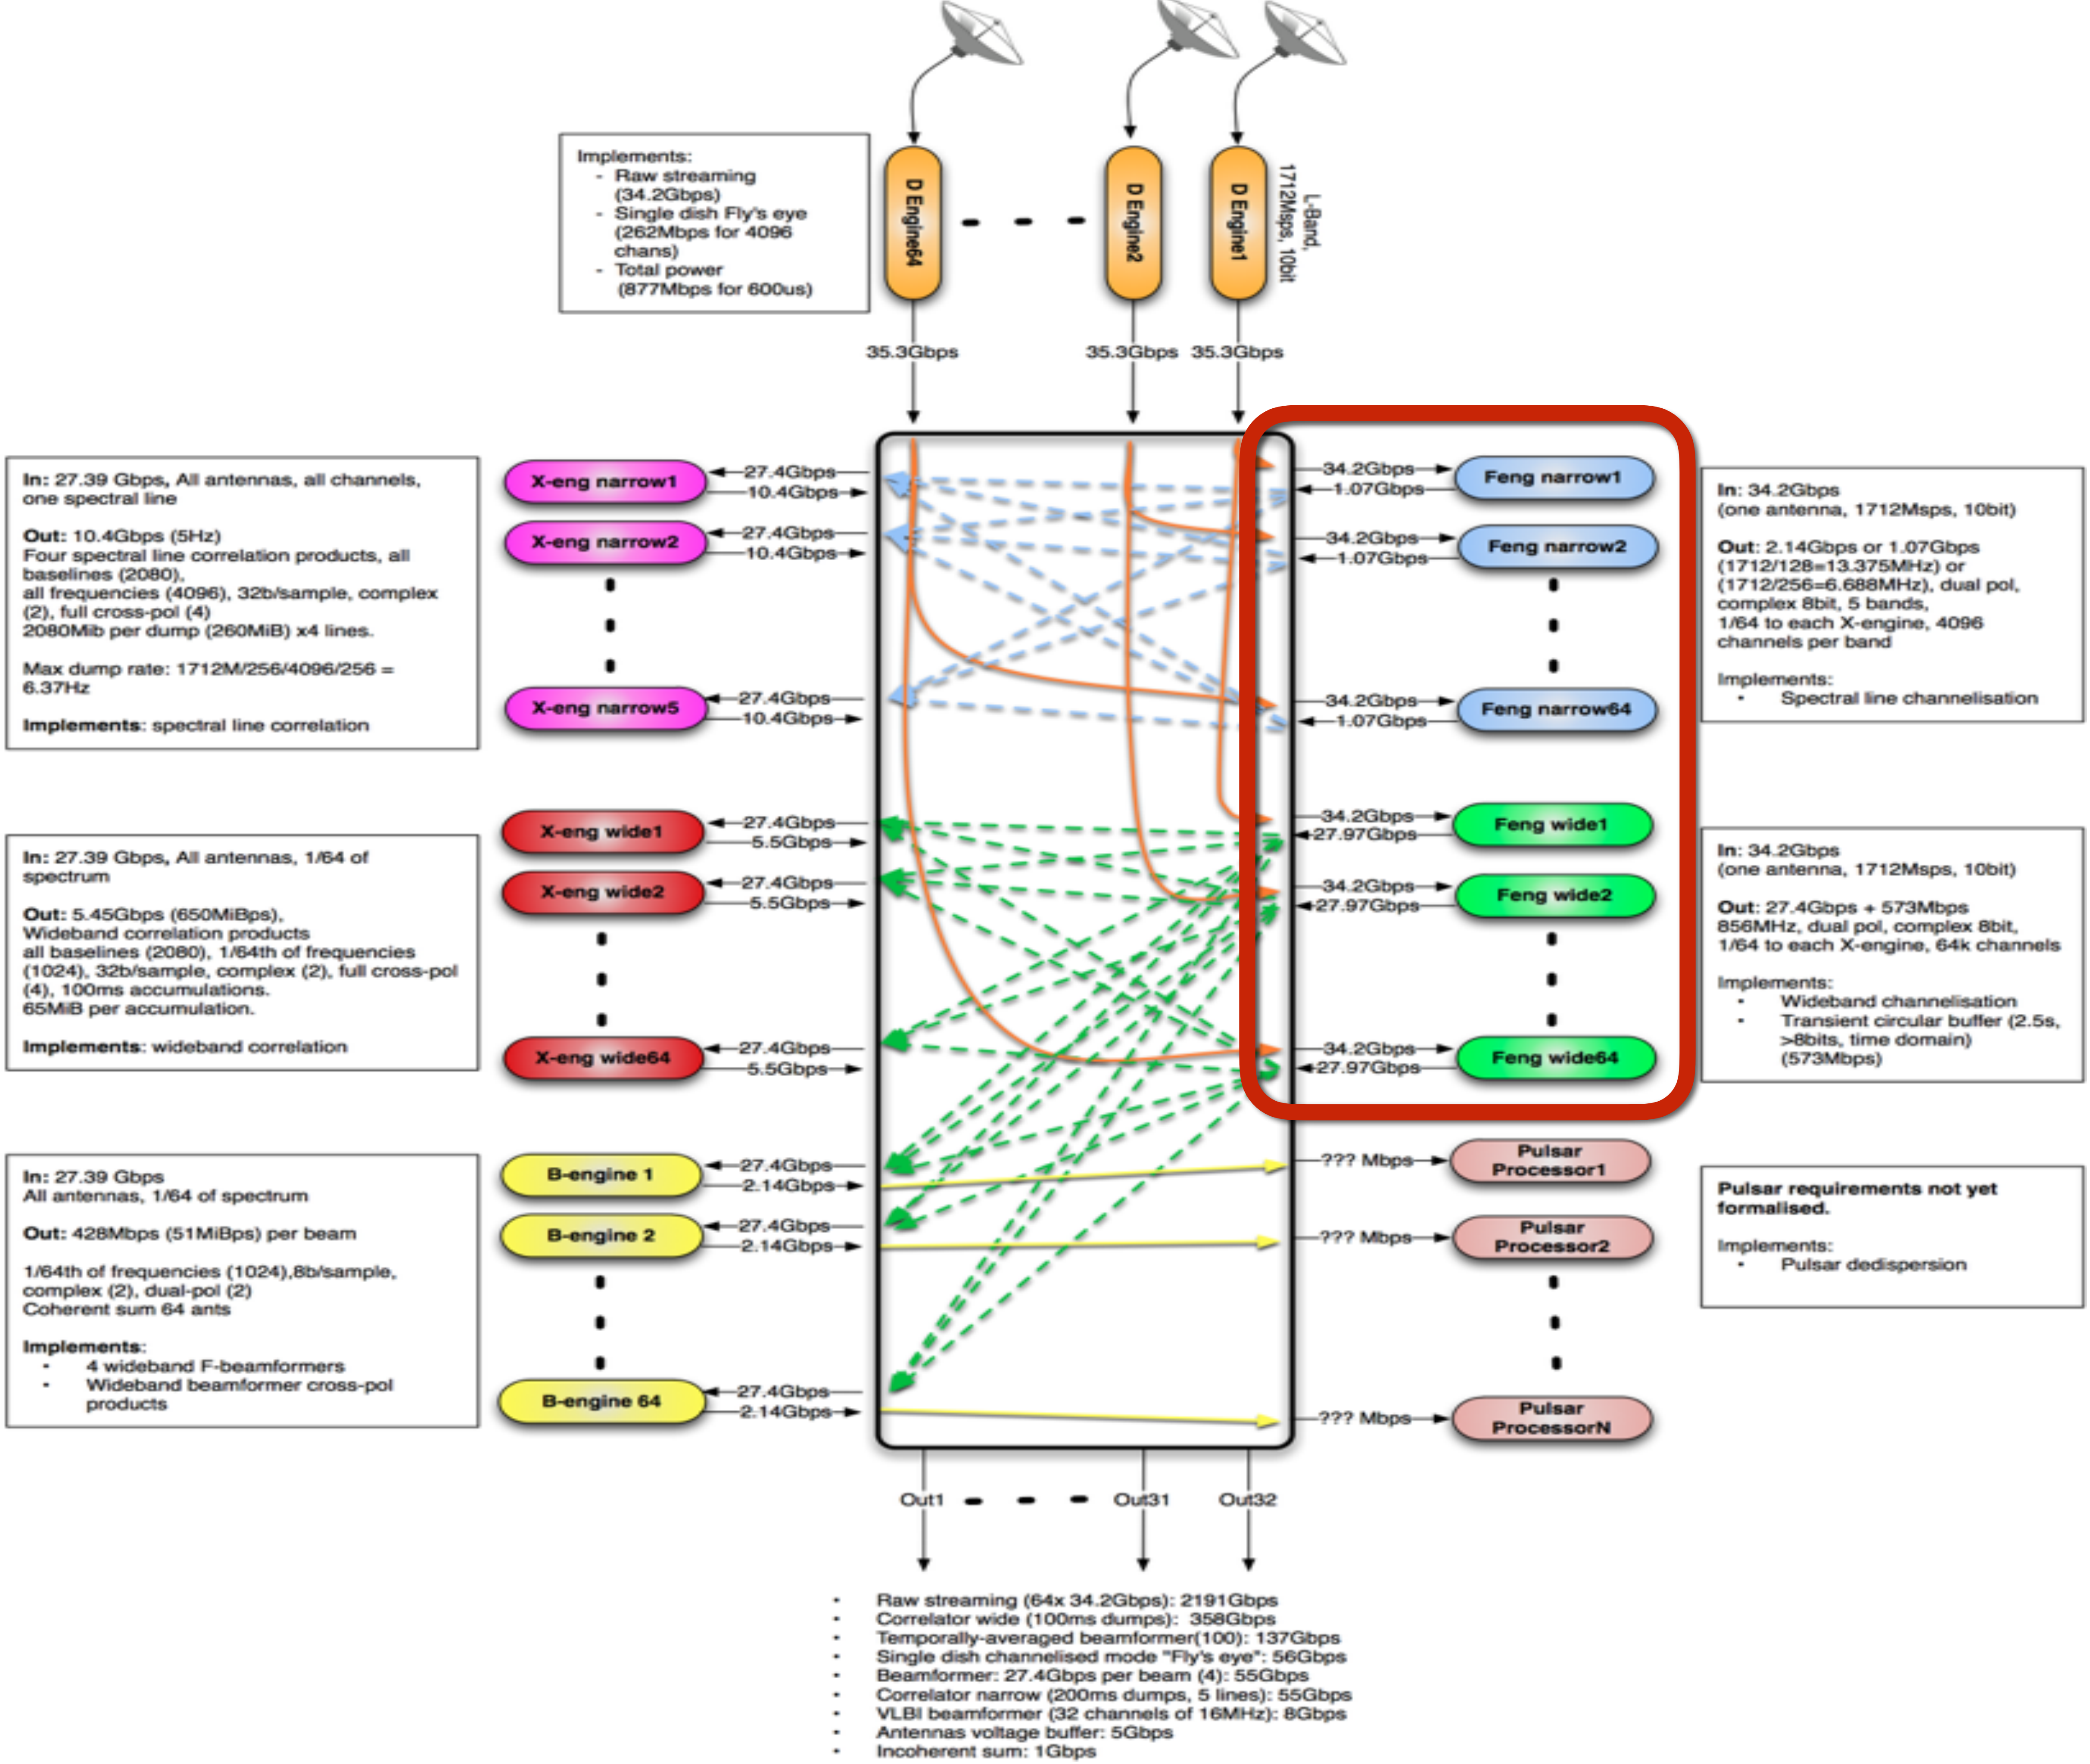

Implements:
- Raw streaming (34.2Gbps)
- Single dish Fly's eye (262Mbps for 4096 chans)
- Total power (877Mbps for 600us)
D Engine64
D Engine2
D Engine1
L-Band, 1712Msps, 10bit
35.3Gbps
35.3Gbps
35.3Gbps
In: 27.39 Gbps, All antennas, all channels, one spectral line
Out: 10.4Gbps (5Hz) Four spectral line correlation products, all baselines (2080), all frequencies (4096), 32b/sample, complex (2), full cross-pol (4) 2080Mib per dump (260MiB) x4 lines.
Max dump rate: 1712M/256/4096/256 = 6.37Hz
Implements: spectral line correlation
X-eng narrow1
X-eng narrow2
X-eng narrow5
27.4Gbps
10.4Gbps
Feng narrow1
Feng narrow2
Feng narrow64
34.2Gbps
1.07Gbps
In: 34.2Gbps (one antenna, 1712Msps, 10bit)
Out: 2.14Gbps or 1.07Gbps (1712/128=13.375MHz) or (1712/256=6.688MHz), dual pol, complex 8bit, 5 bands, 1/64 to each X-engine, 4096 channels per band
Implements:
• Spectral line channelisation
In: 27.39 Gbps, All antennas, 1/64 of spectrum
Out: 5.45Gbps (650MiBps), Wideband correlation products all baselines (2080), 1/64th of frequencies (1024), 32b/sample, complex (2), full cross-pol (4), 100ms accumulations. 65MiB per accumulation.
Implements: wideband correlation
X-eng wide1
X-eng wide2
X-eng wide64
27.4Gbps
5.5Gbps
Feng wide1
Feng wide2
Feng wide64
34.2Gbps
27.97Gbps
In: 34.2Gbps (one antenna, 1712Msps, 10bit)
Out: 27.4Gbps + 573Mbps 856MHz, dual pol, complex 8bit, 1/64 to each X-engine, 64k channels
Implements:
• Wideband channelisation
• Transient circular buffer (2.5s, >8bits, time domain) (573Mbps)
In: 27.39 Gbps All antennas, 1/64 of spectrum
Out: 428Mbps (51MiBps) per beam
1/64th of frequencies (1024),8b/sample, complex (2), dual-pol (2) Coherent sum 64 ants
Implements:
• 4 wideband F-beamformers
• Wideband beamformer cross-pol products
B-engine 1
B-engine 2
B-engine 64
27.4Gbps
2.14Gbps
Pulsar Processor1
Pulsar Processor2
Pulsar ProcessorN
??? Mbps
Pulsar requirements not yet formalised.
Implements:
• Pulsar dedispersion
Out1
Out31
Out32
• Raw streaming (64x 34.2Gbps): 2191Gbps
• Correlator wide (100ms dumps): 358Gbps
• Temporally-averaged beamformer(100): 137Gbps
• Single dish channelised mode "Fly's eye": 56Gbps
• Beamformer: 27.4Gbps per beam (4): 55Gbps
• Correlator narrow (200ms dumps, 5 lines): 55Gbps
• VLBI beamformer (32 channels of 16MHz): 8Gbps
• Antennas voltage buffer: 5Gbps
• Incoherent sum: 1Gbps

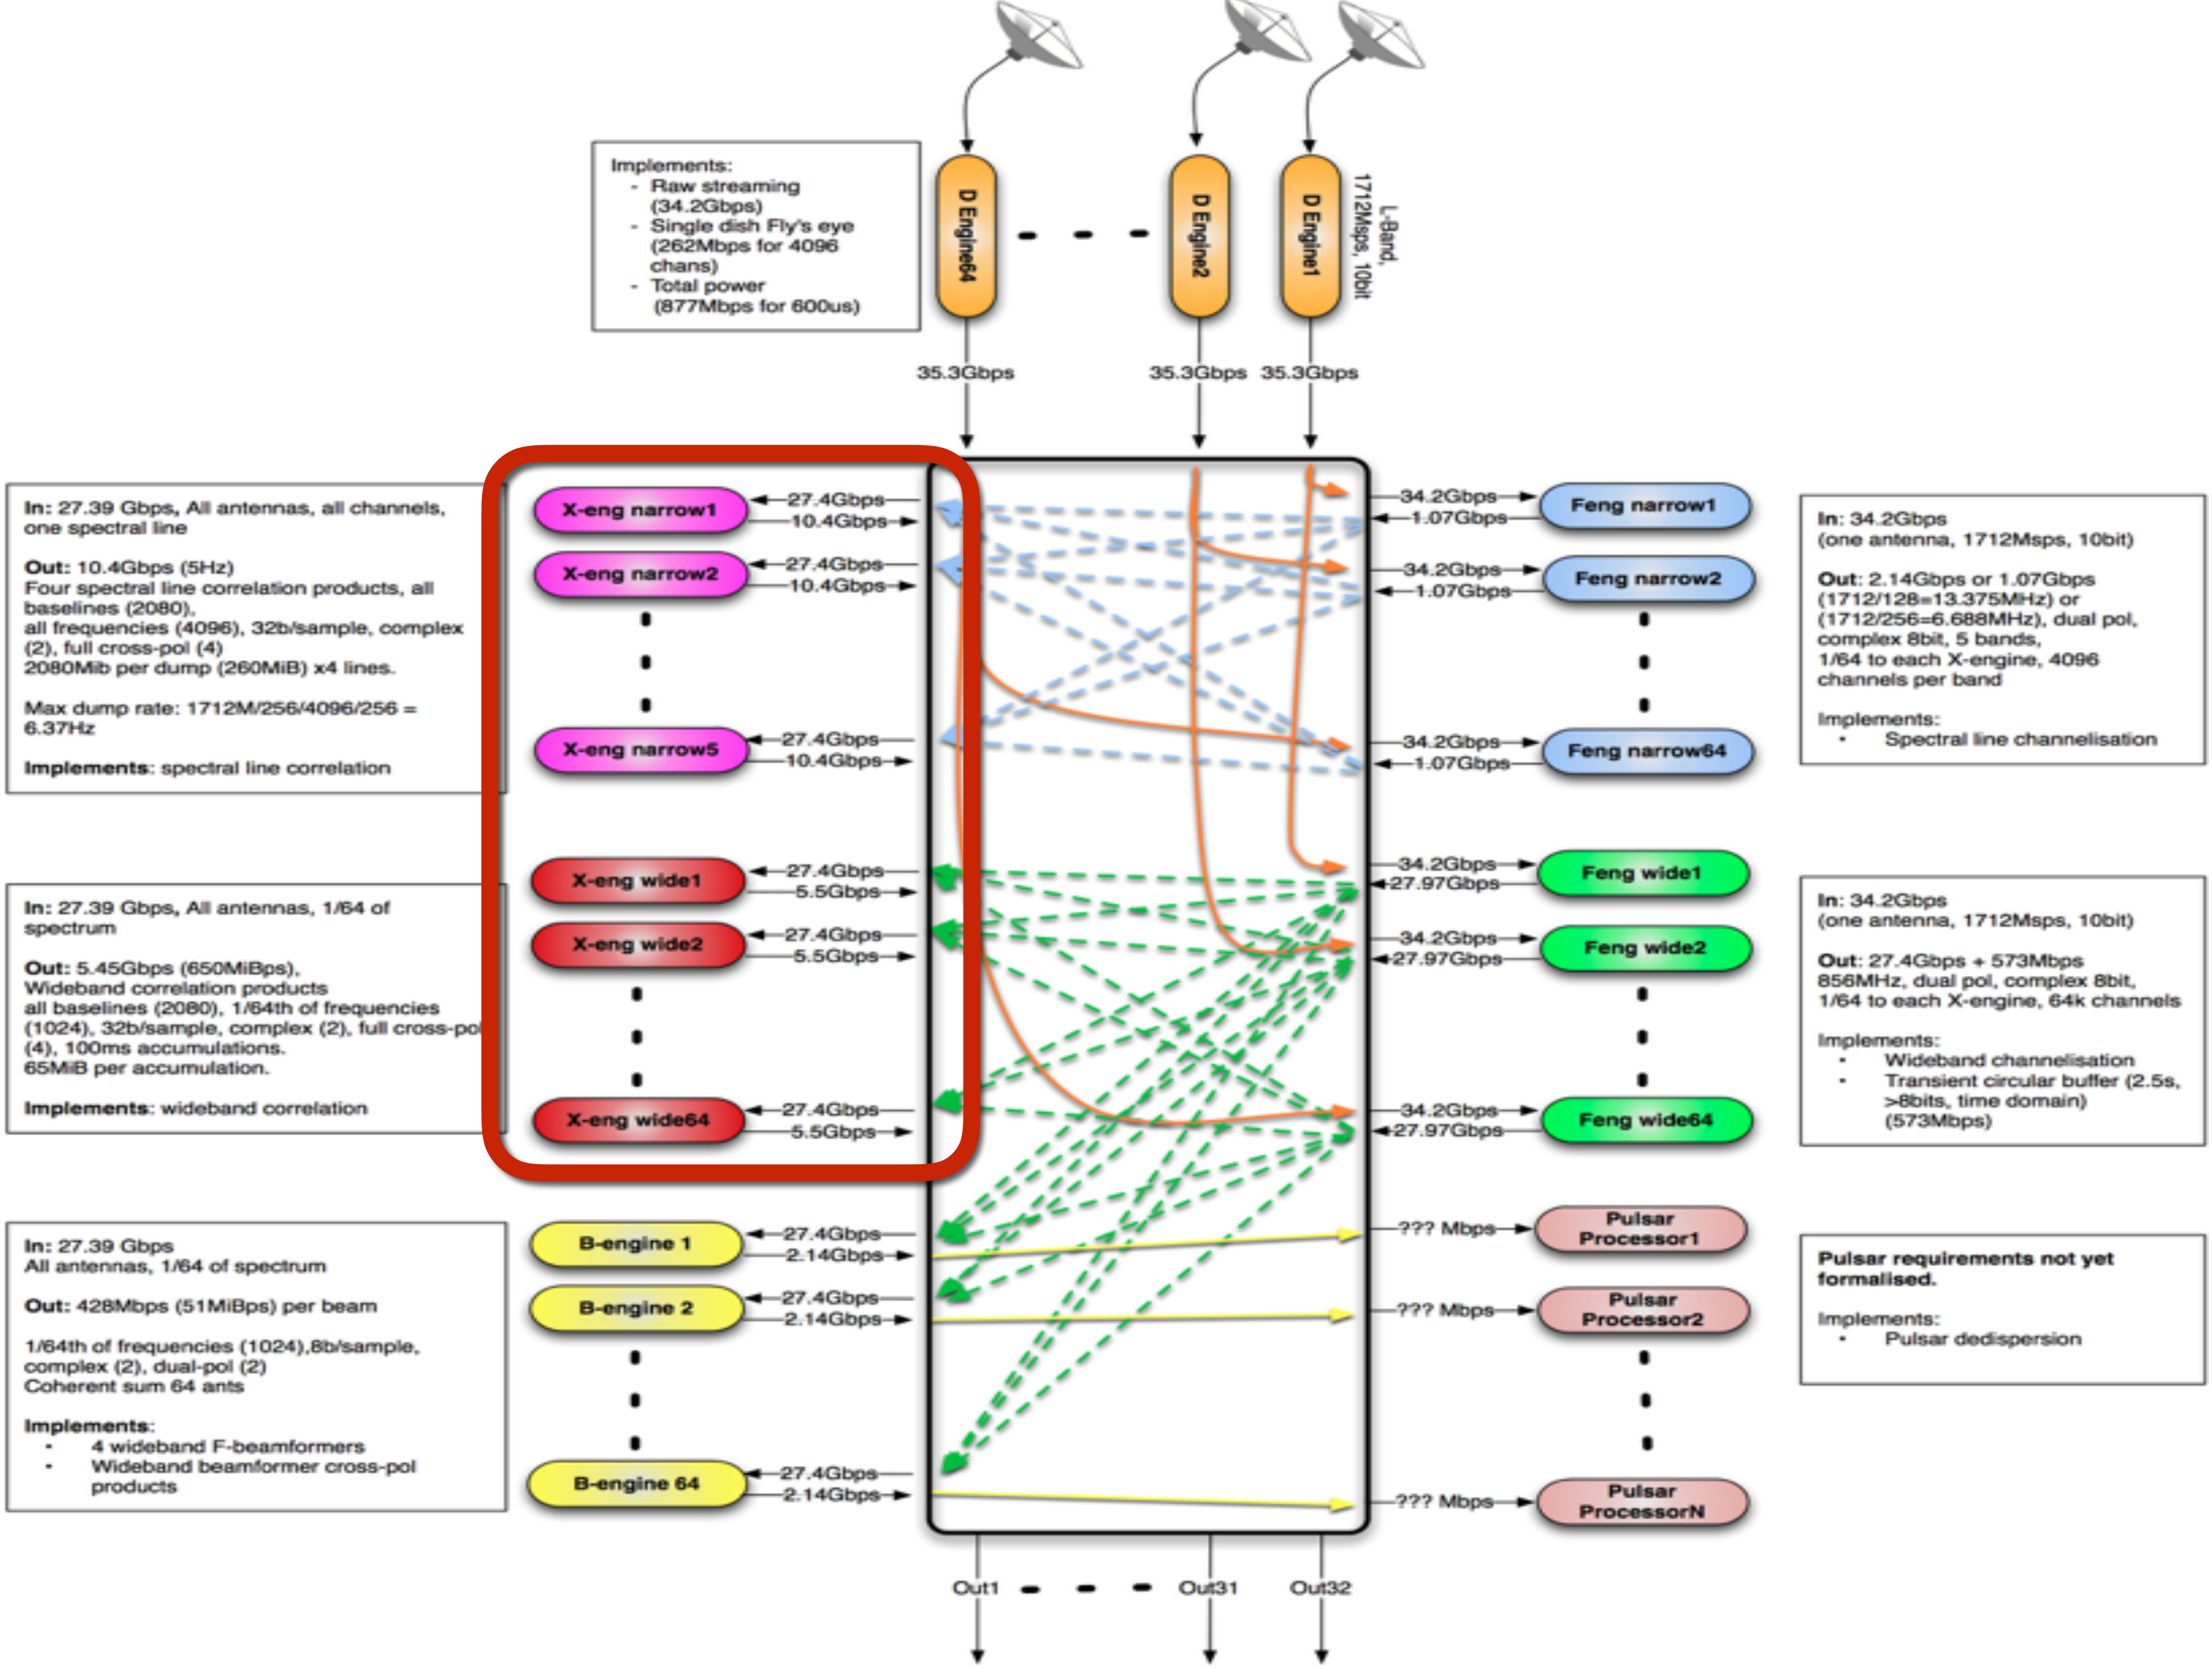

- Raw streaming (64x 34.2Gbps): 2191Gbps
- Correlator wide (100ms dumps): 358Gbps
- Temporally-averaged beamformer(100): 137Gbps
- Single dish channelised mode "Fly's eye": 56Gbps
- Beamformer: 27.4Gbps per beam (4): 55Gbps
- Correlator narrow (200ms dumps, 5 lines): 55Gbps
- VLBI beamformer (32 channels of 16MHz): 8Gbps
- Antennas voltage buffer: 5Gbps
- Incoherent sum: 1Gbps

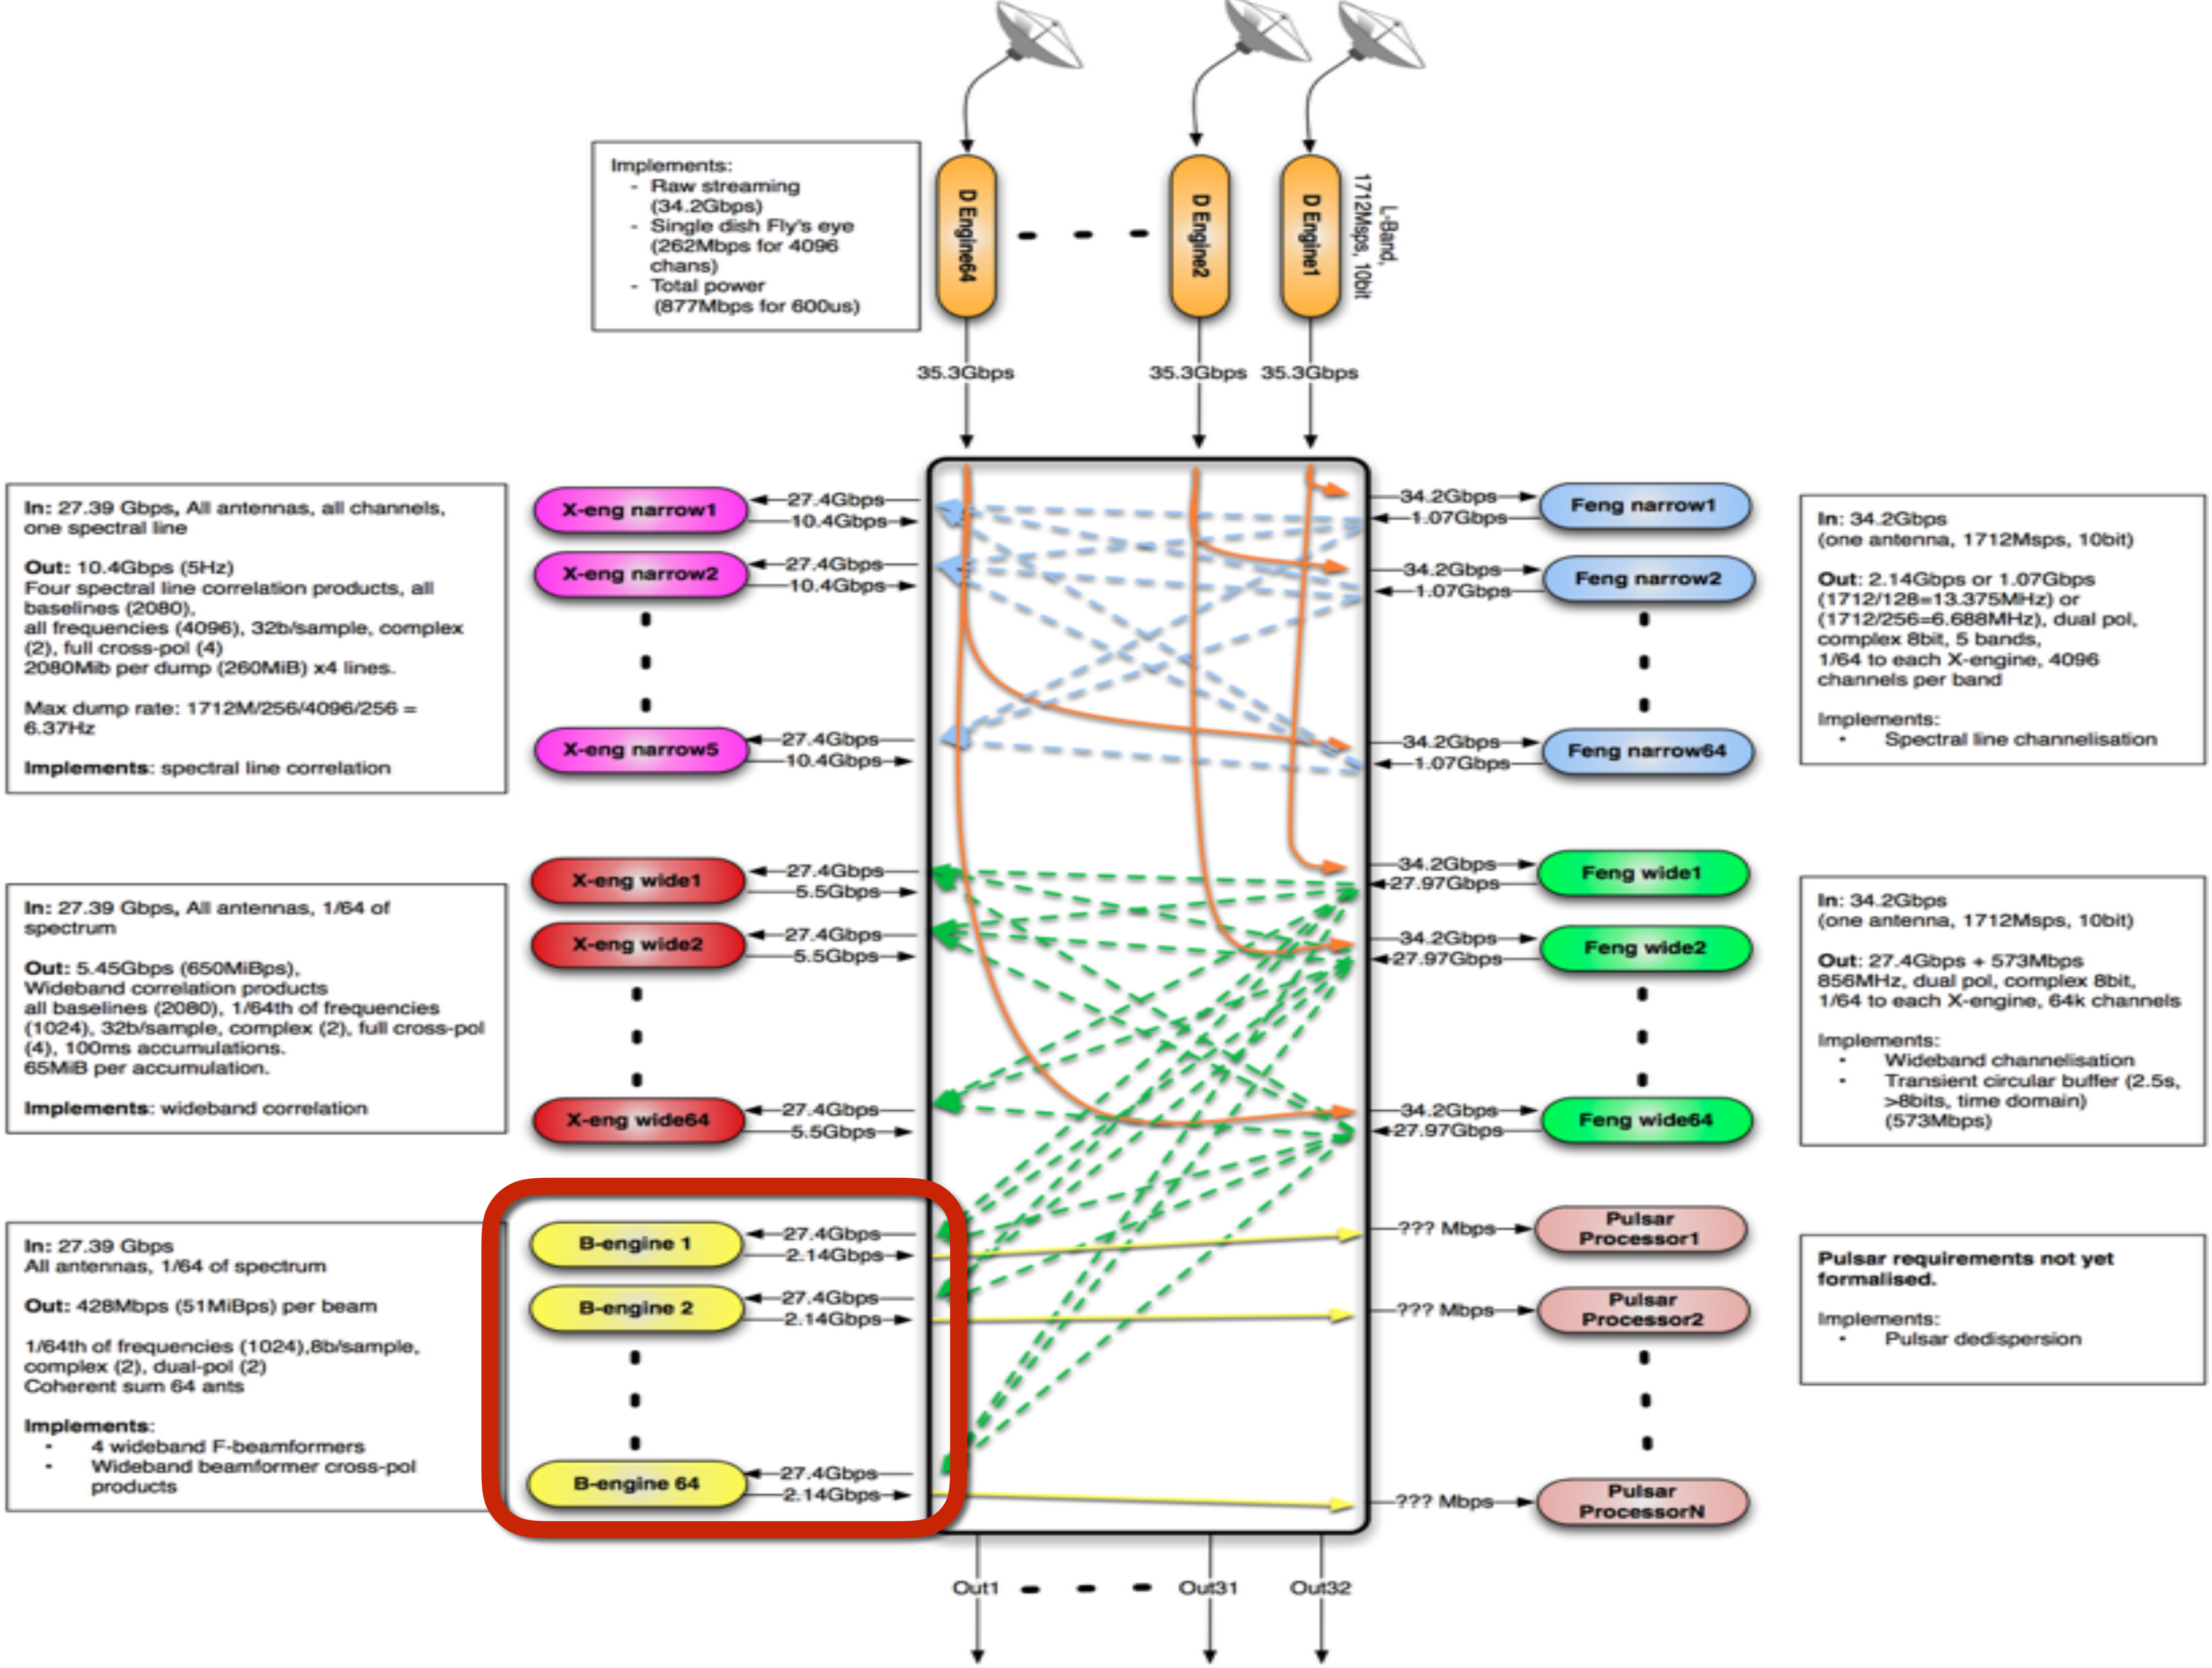

- Raw streaming (64x 34.2Gbps): 2191Gbps
- Correlator wide (100ms dumps): 358Gbps
- Temporally-averaged beamformer(100): 137Gbps
- Single dish channelised mode "Fly's eye": 56Gbps
- Beamformer: 27.4Gbps per beam (4): 55Gbps
- Correlator narrow (200ms dumps, 5 lines): 55Gbps
- VLBI beamformer (32 channels of 16MHz): 8Gbps
- Antennas voltage buffer: 5Gbps
- Incoherent sum: 1Gbps

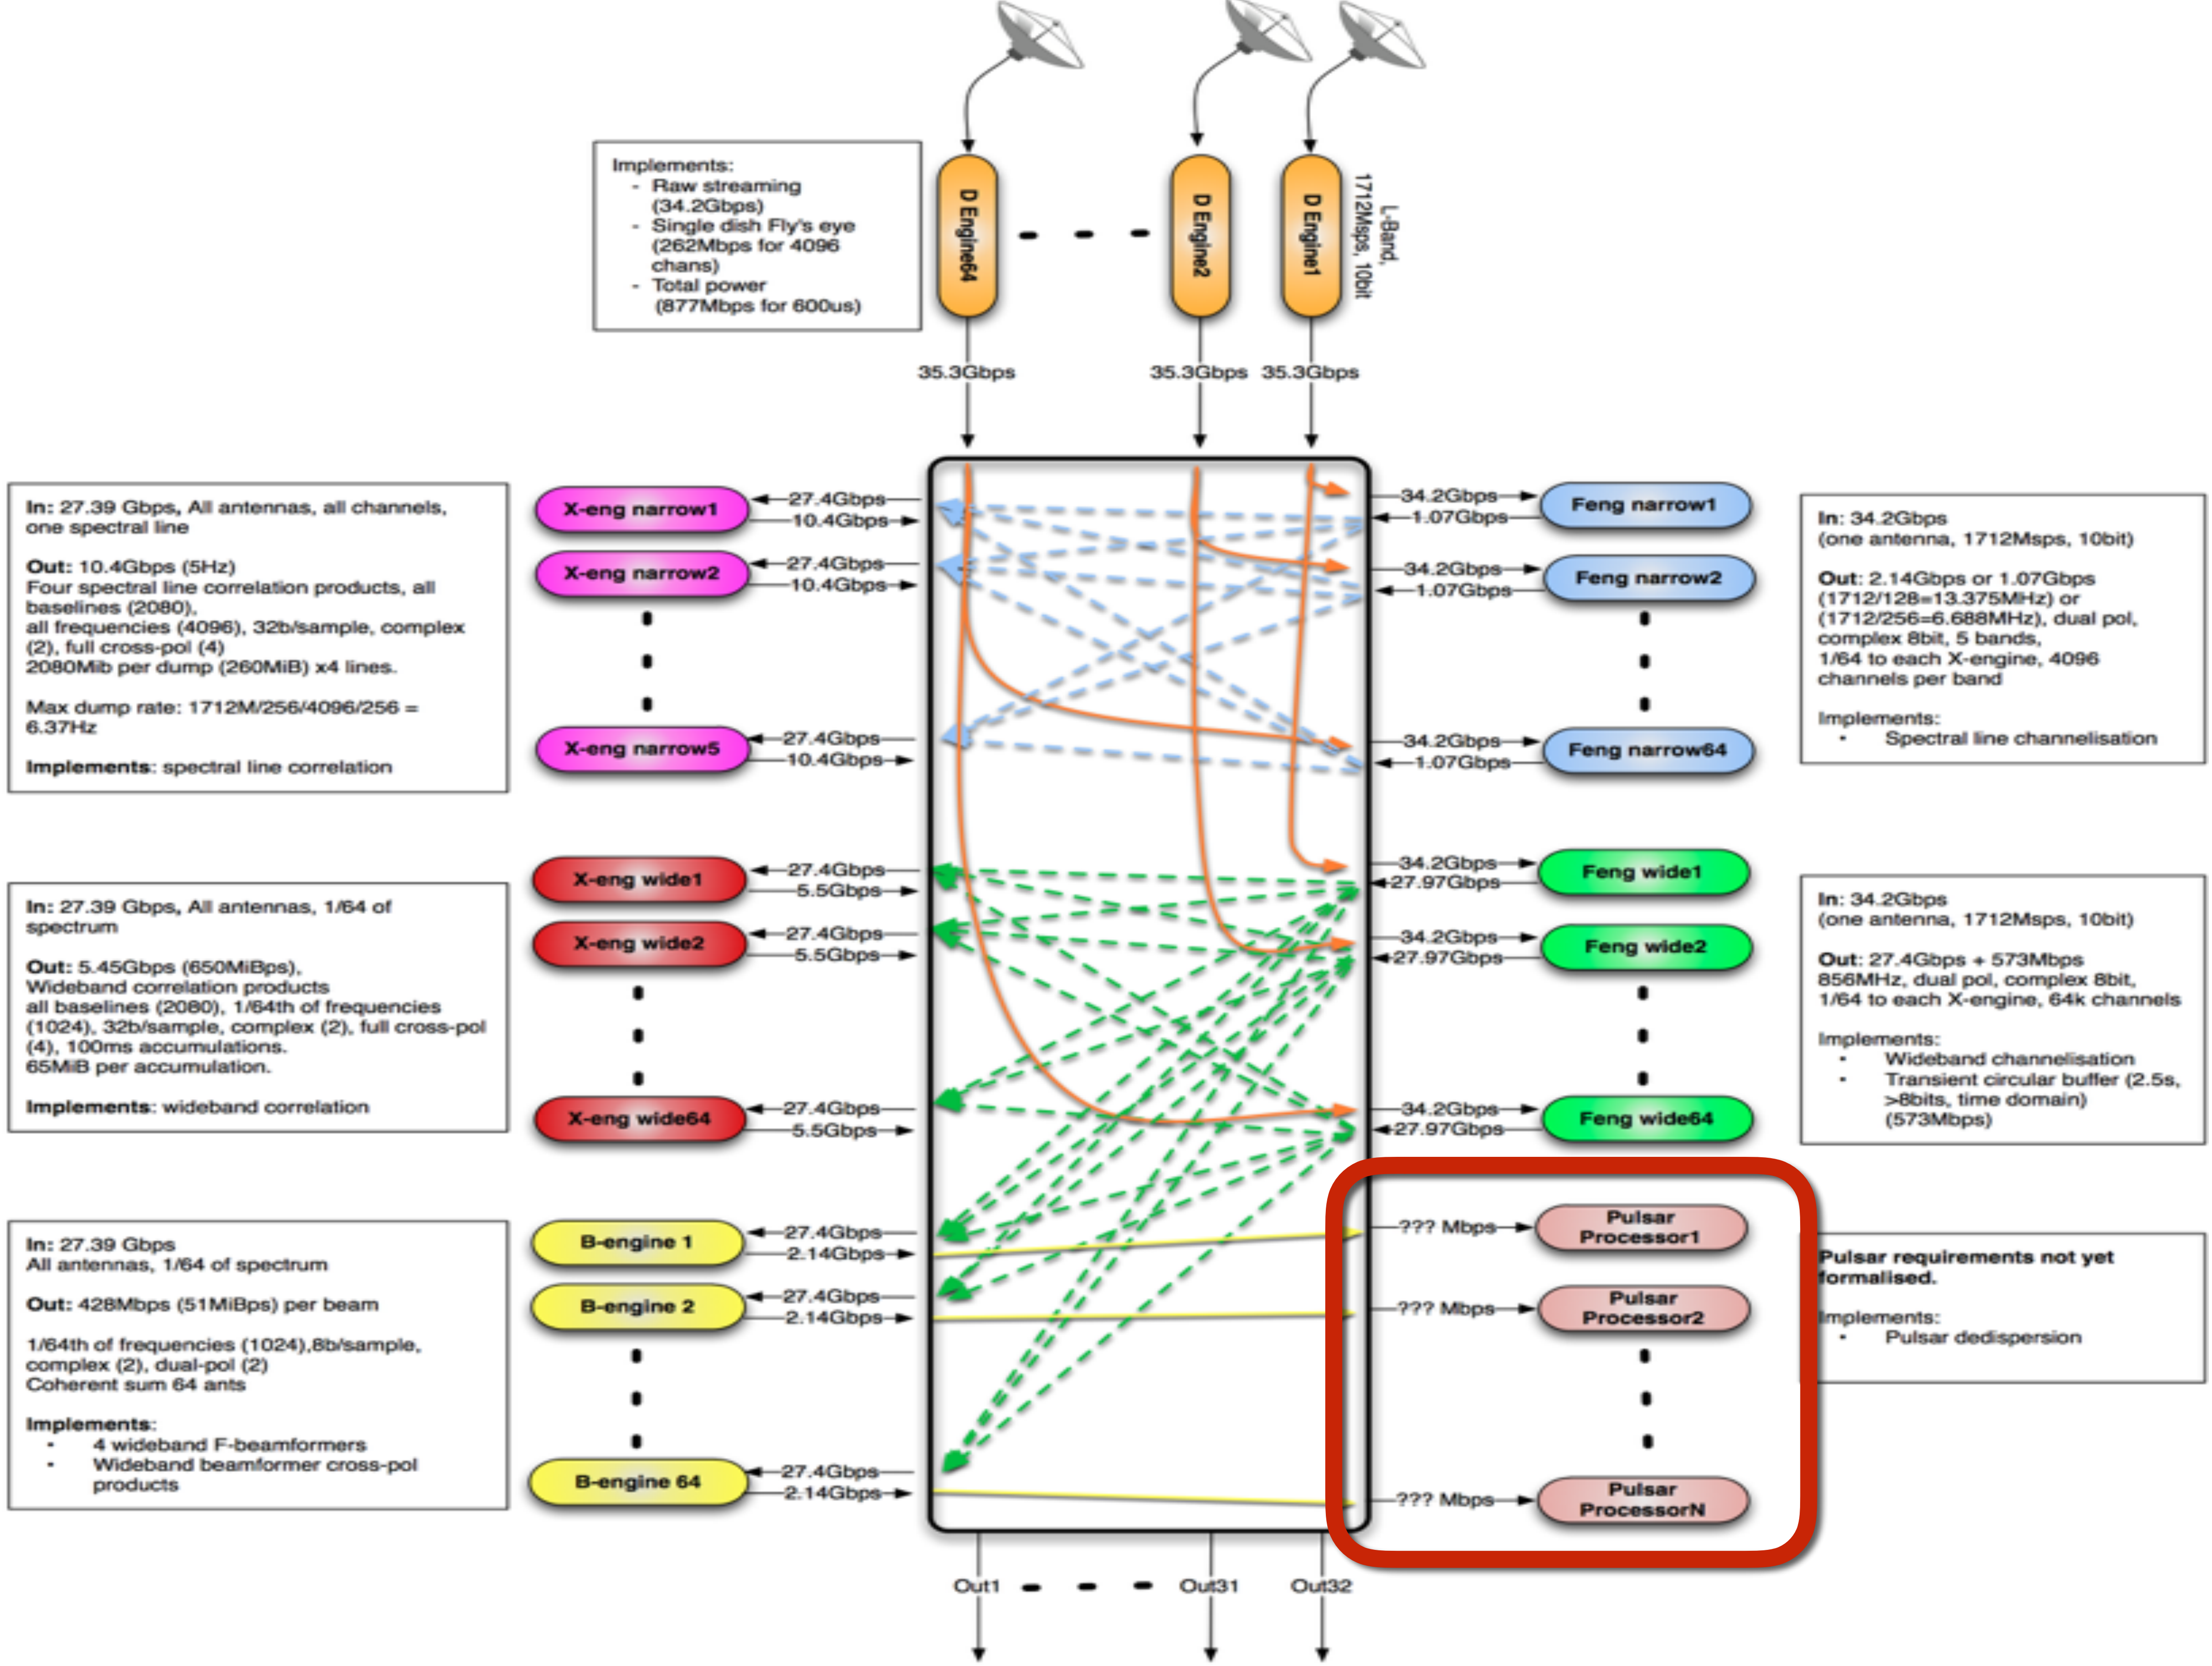

- Raw streaming (64x 34.2Gbps): 2191Gbps
- Correlator wide (100ms dumps): 358Gbps
- Temporally-averaged beamformer(100): 137Gbps
- Single dish channelised mode "Fly's eye": 56Gbps
- Beamformer: 27.4Gbps per beam (4): 55Gbps
- Correlator narrow (200ms dumps, 5 lines): 55Gbps
- VLBI beamformer (32 channels of 16MHz): 8Gbps
- Antennas voltage buffer: 5Gbps
- Incoherent sum: 1Gbps

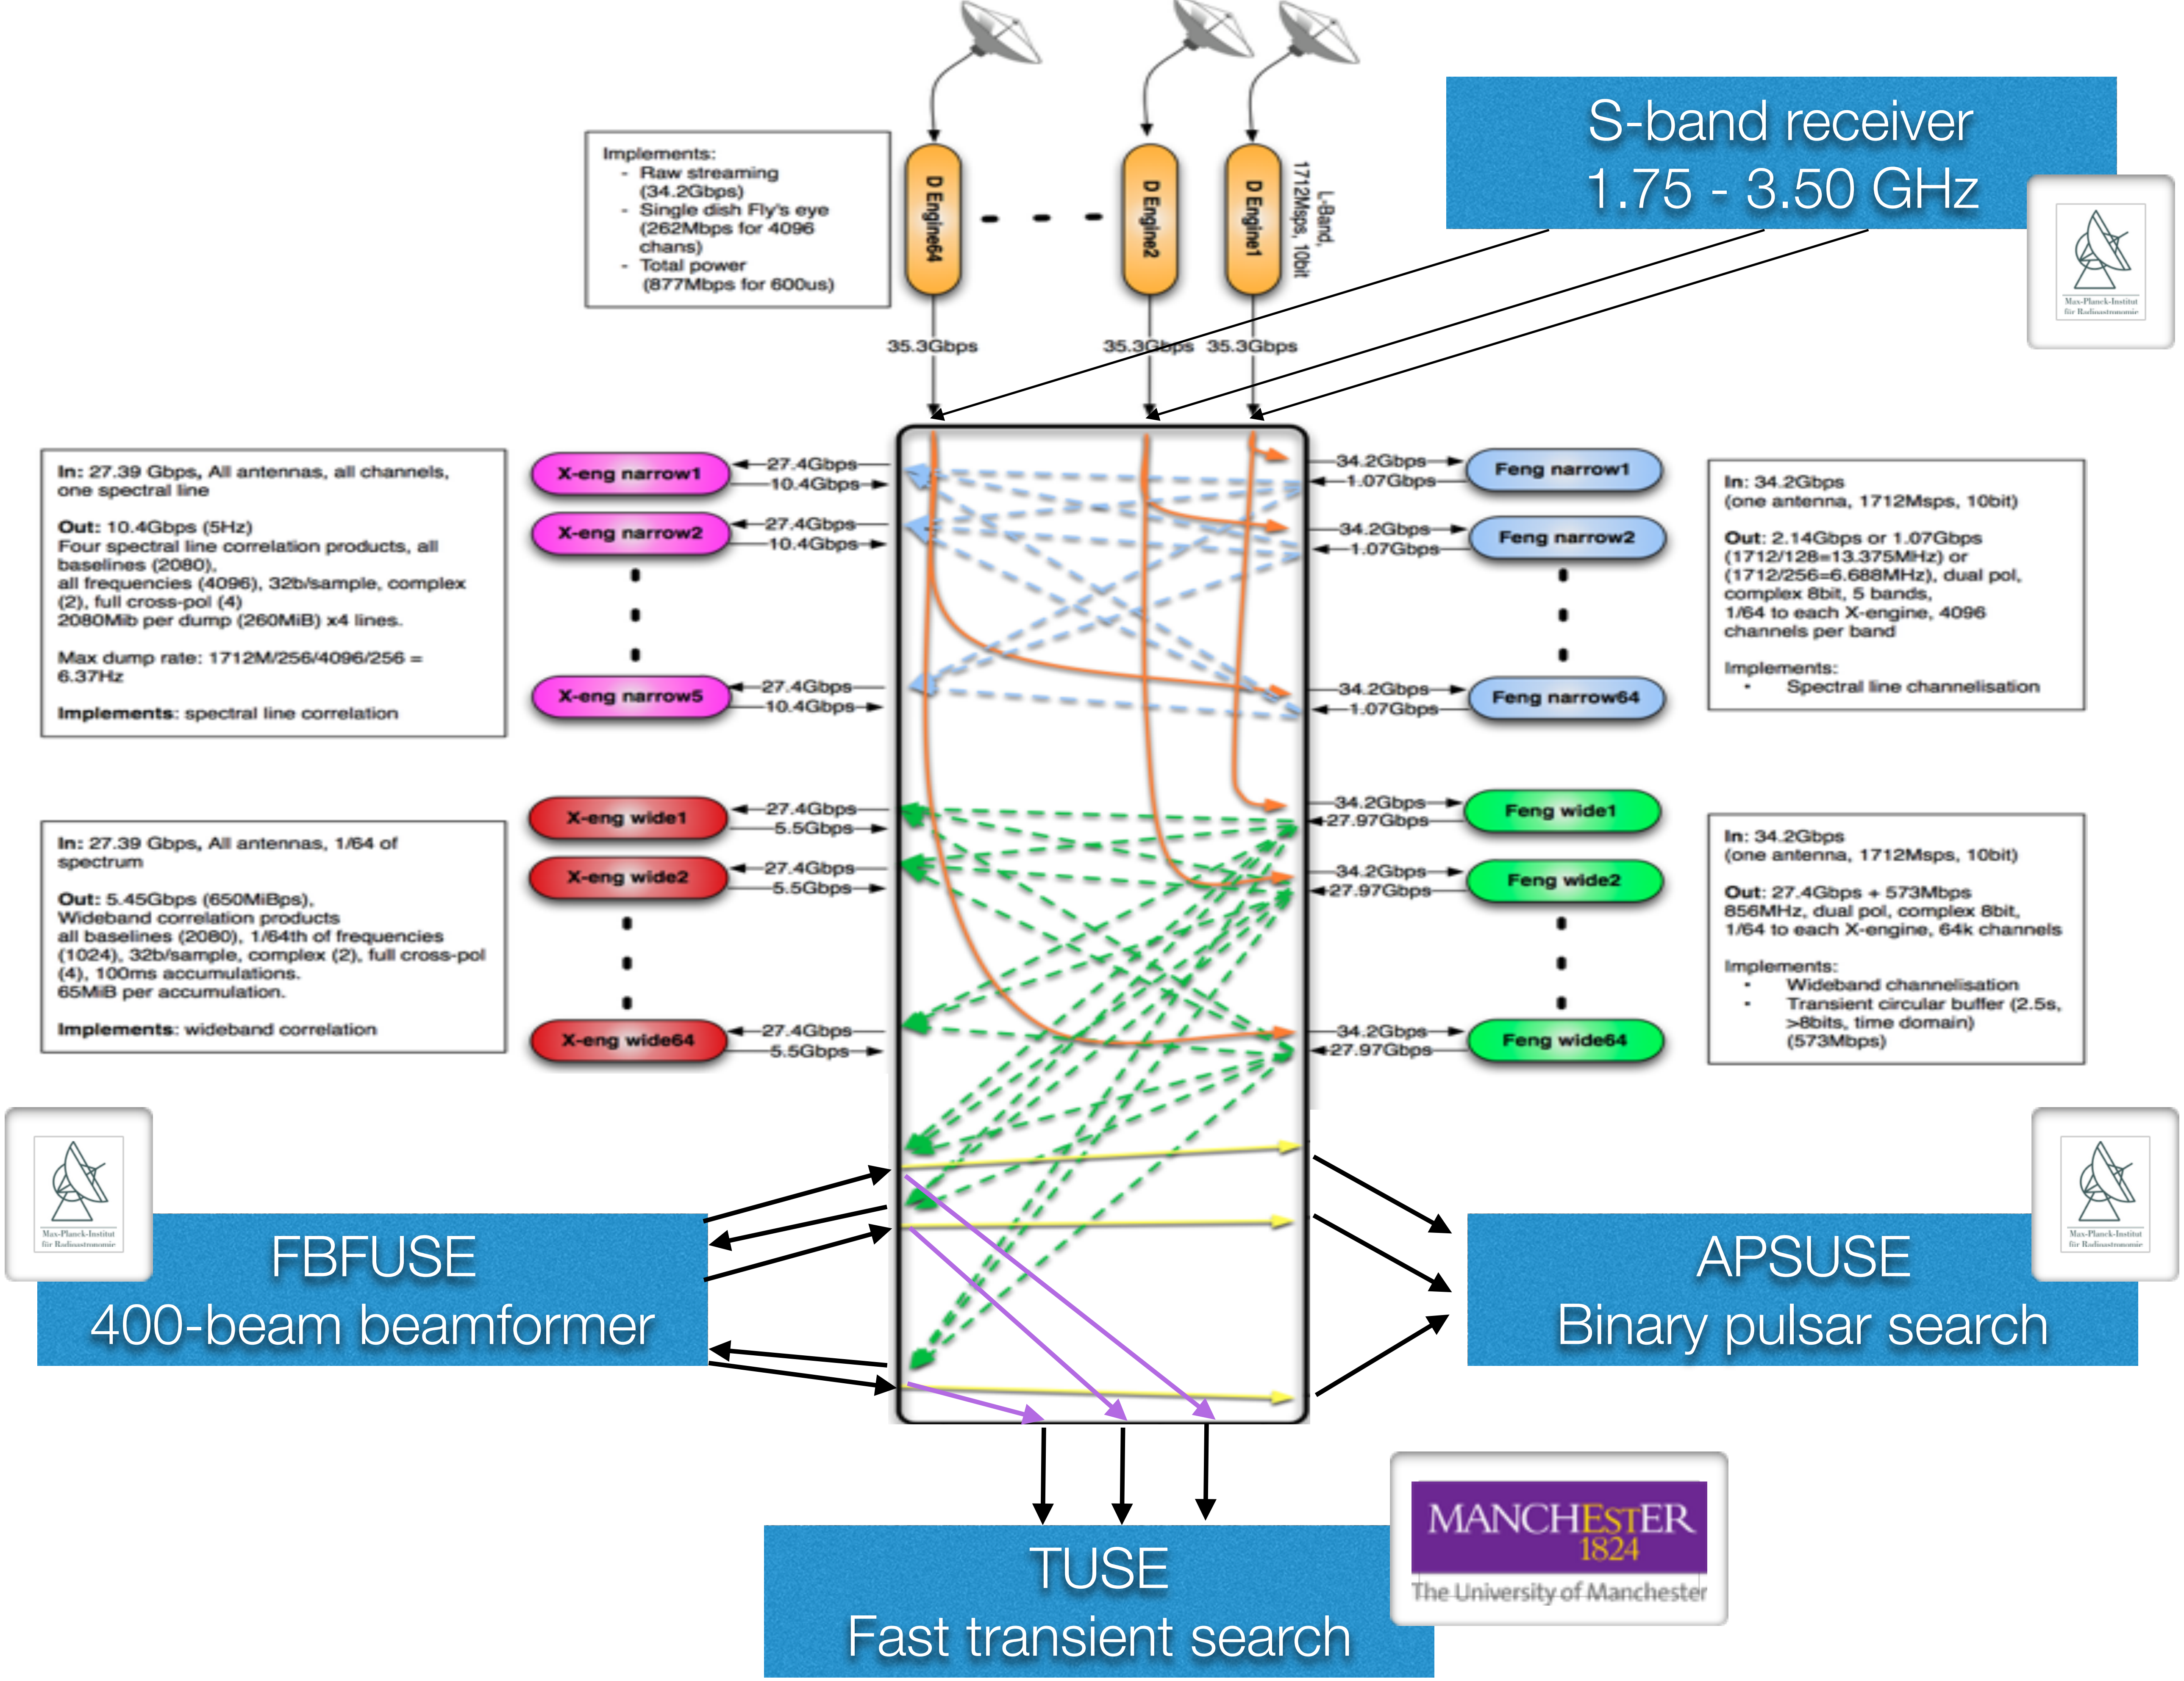

S-band receiver
1.75 - 3.50 GHz
Implements:
- Raw streaming (34.2Gbps)
- Single dish Fly's eye (262Mbps for 4096 chans)
- Total power (877Mbps for 600us)
D Engine64
D Engine2
D Engine1
L-Band, 1712Msps, 10bit
35.3Gbps
35.3Gbps
35.3Gbps
In: 27.39 Gbps, All antennas, all channels, one spectral line
Out: 10.4Gbps (5Hz)
Four spectral line correlation products, all baselines (2080), all frequencies (4096), 32b/sample, complex (2), full cross-pol (4)
2080Mib per dump (260MiB) x4 lines.
Max dump rate: 1712M/256/4096/256 = 6.37Hz
Implements: spectral line correlation
X-eng narrow1
X-eng narrow2
X-eng narrow5
27.4Gbps
10.4Gbps
Feng narrow1
Feng narrow2
Feng narrow64
34.2Gbps
1.07Gbps
In: 34.2Gbps (one antenna, 1712Msps, 10bit)
Out: 2.14Gbps or 1.07Gbps (1712/128=13.375MHz) or (1712/256=6.688MHz), dual pol, complex 8bit, 5 bands, 1/64 to each X-engine, 4096 channels per band
Implements:
- Spectral line channelisation
In: 27.39 Gbps, All antennas, 1/64 of spectrum
Out: 5.45Gbps (650MiBps), Wideband correlation products all baselines (2080), 1/64th of frequencies (1024), 32b/sample, complex (2), full cross-pol (4), 100ms accumulations.
65MiB per accumulation.
Implements: wideband correlation
X-eng wide1
X-eng wide2
X-eng wide64
27.4Gbps
5.5Gbps
Feng wide1
Feng wide2
Feng wide64
34.2Gbps
27.97Gbps
In: 34.2Gbps (one antenna, 1712Msps, 10bit)
Out: 27.4Gbps + 573Mbps 856MHz, dual pol, complex 8bit, 1/64 to each X-engine, 64k channels
Implements:
- Wideband channelisation
- Transient circular buffer (2.5s, >8bits, time domain) (573Mbps)
FBFUSE
400-beam beamformer
APSUSE
Binary pulsar search
TUSE
Fast transient search
MANCHESTER 1824
The University of Manchester

Science goals

- Fast Radio Bursts:
  - Emission properties
  - Localisation
  - Multimessenger studies: Optical, gamma-ray, neutrino, etc.
- Binary pulsars:
  - Test GR in the quasi-stationary strong-field regime.
  - Testing SEP violation
  - Testing alternate theories of gravity (e.g. TeVeS)

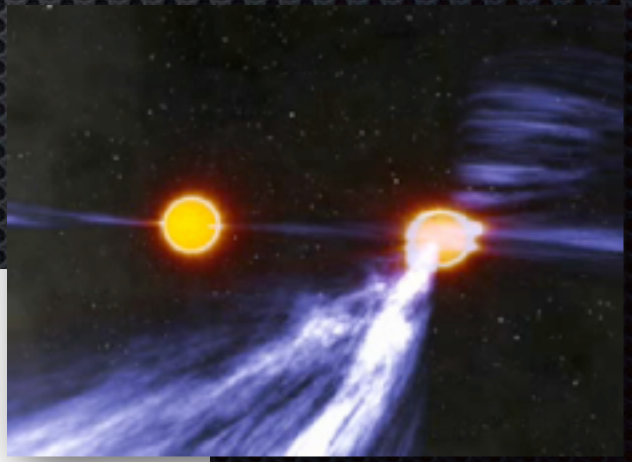

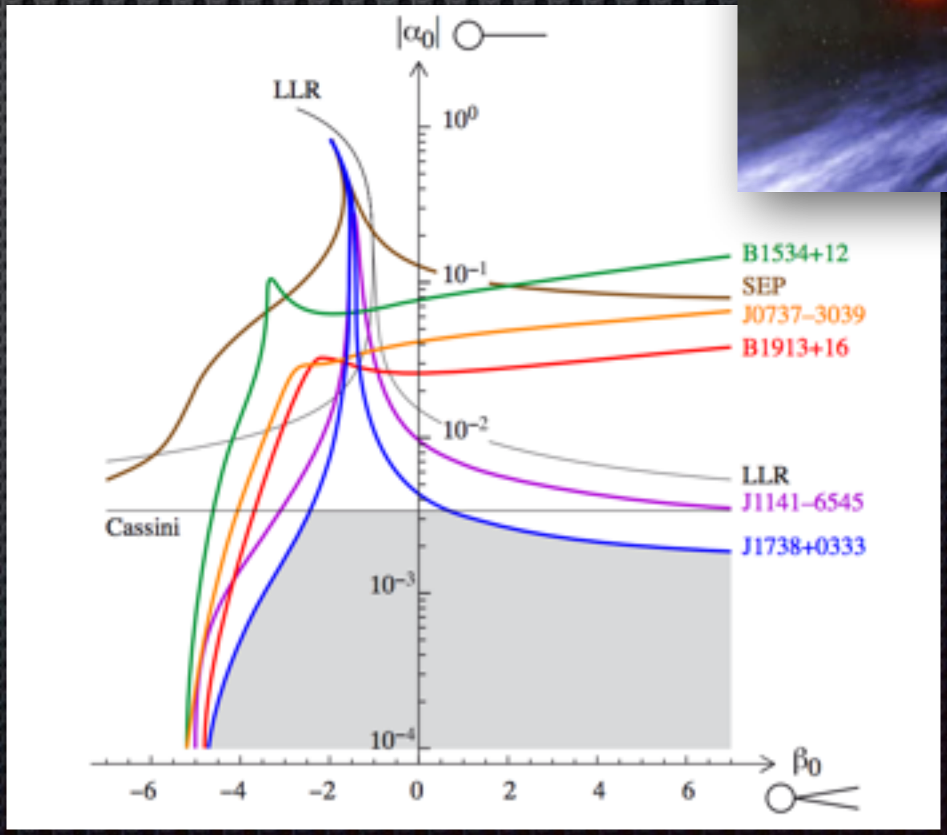

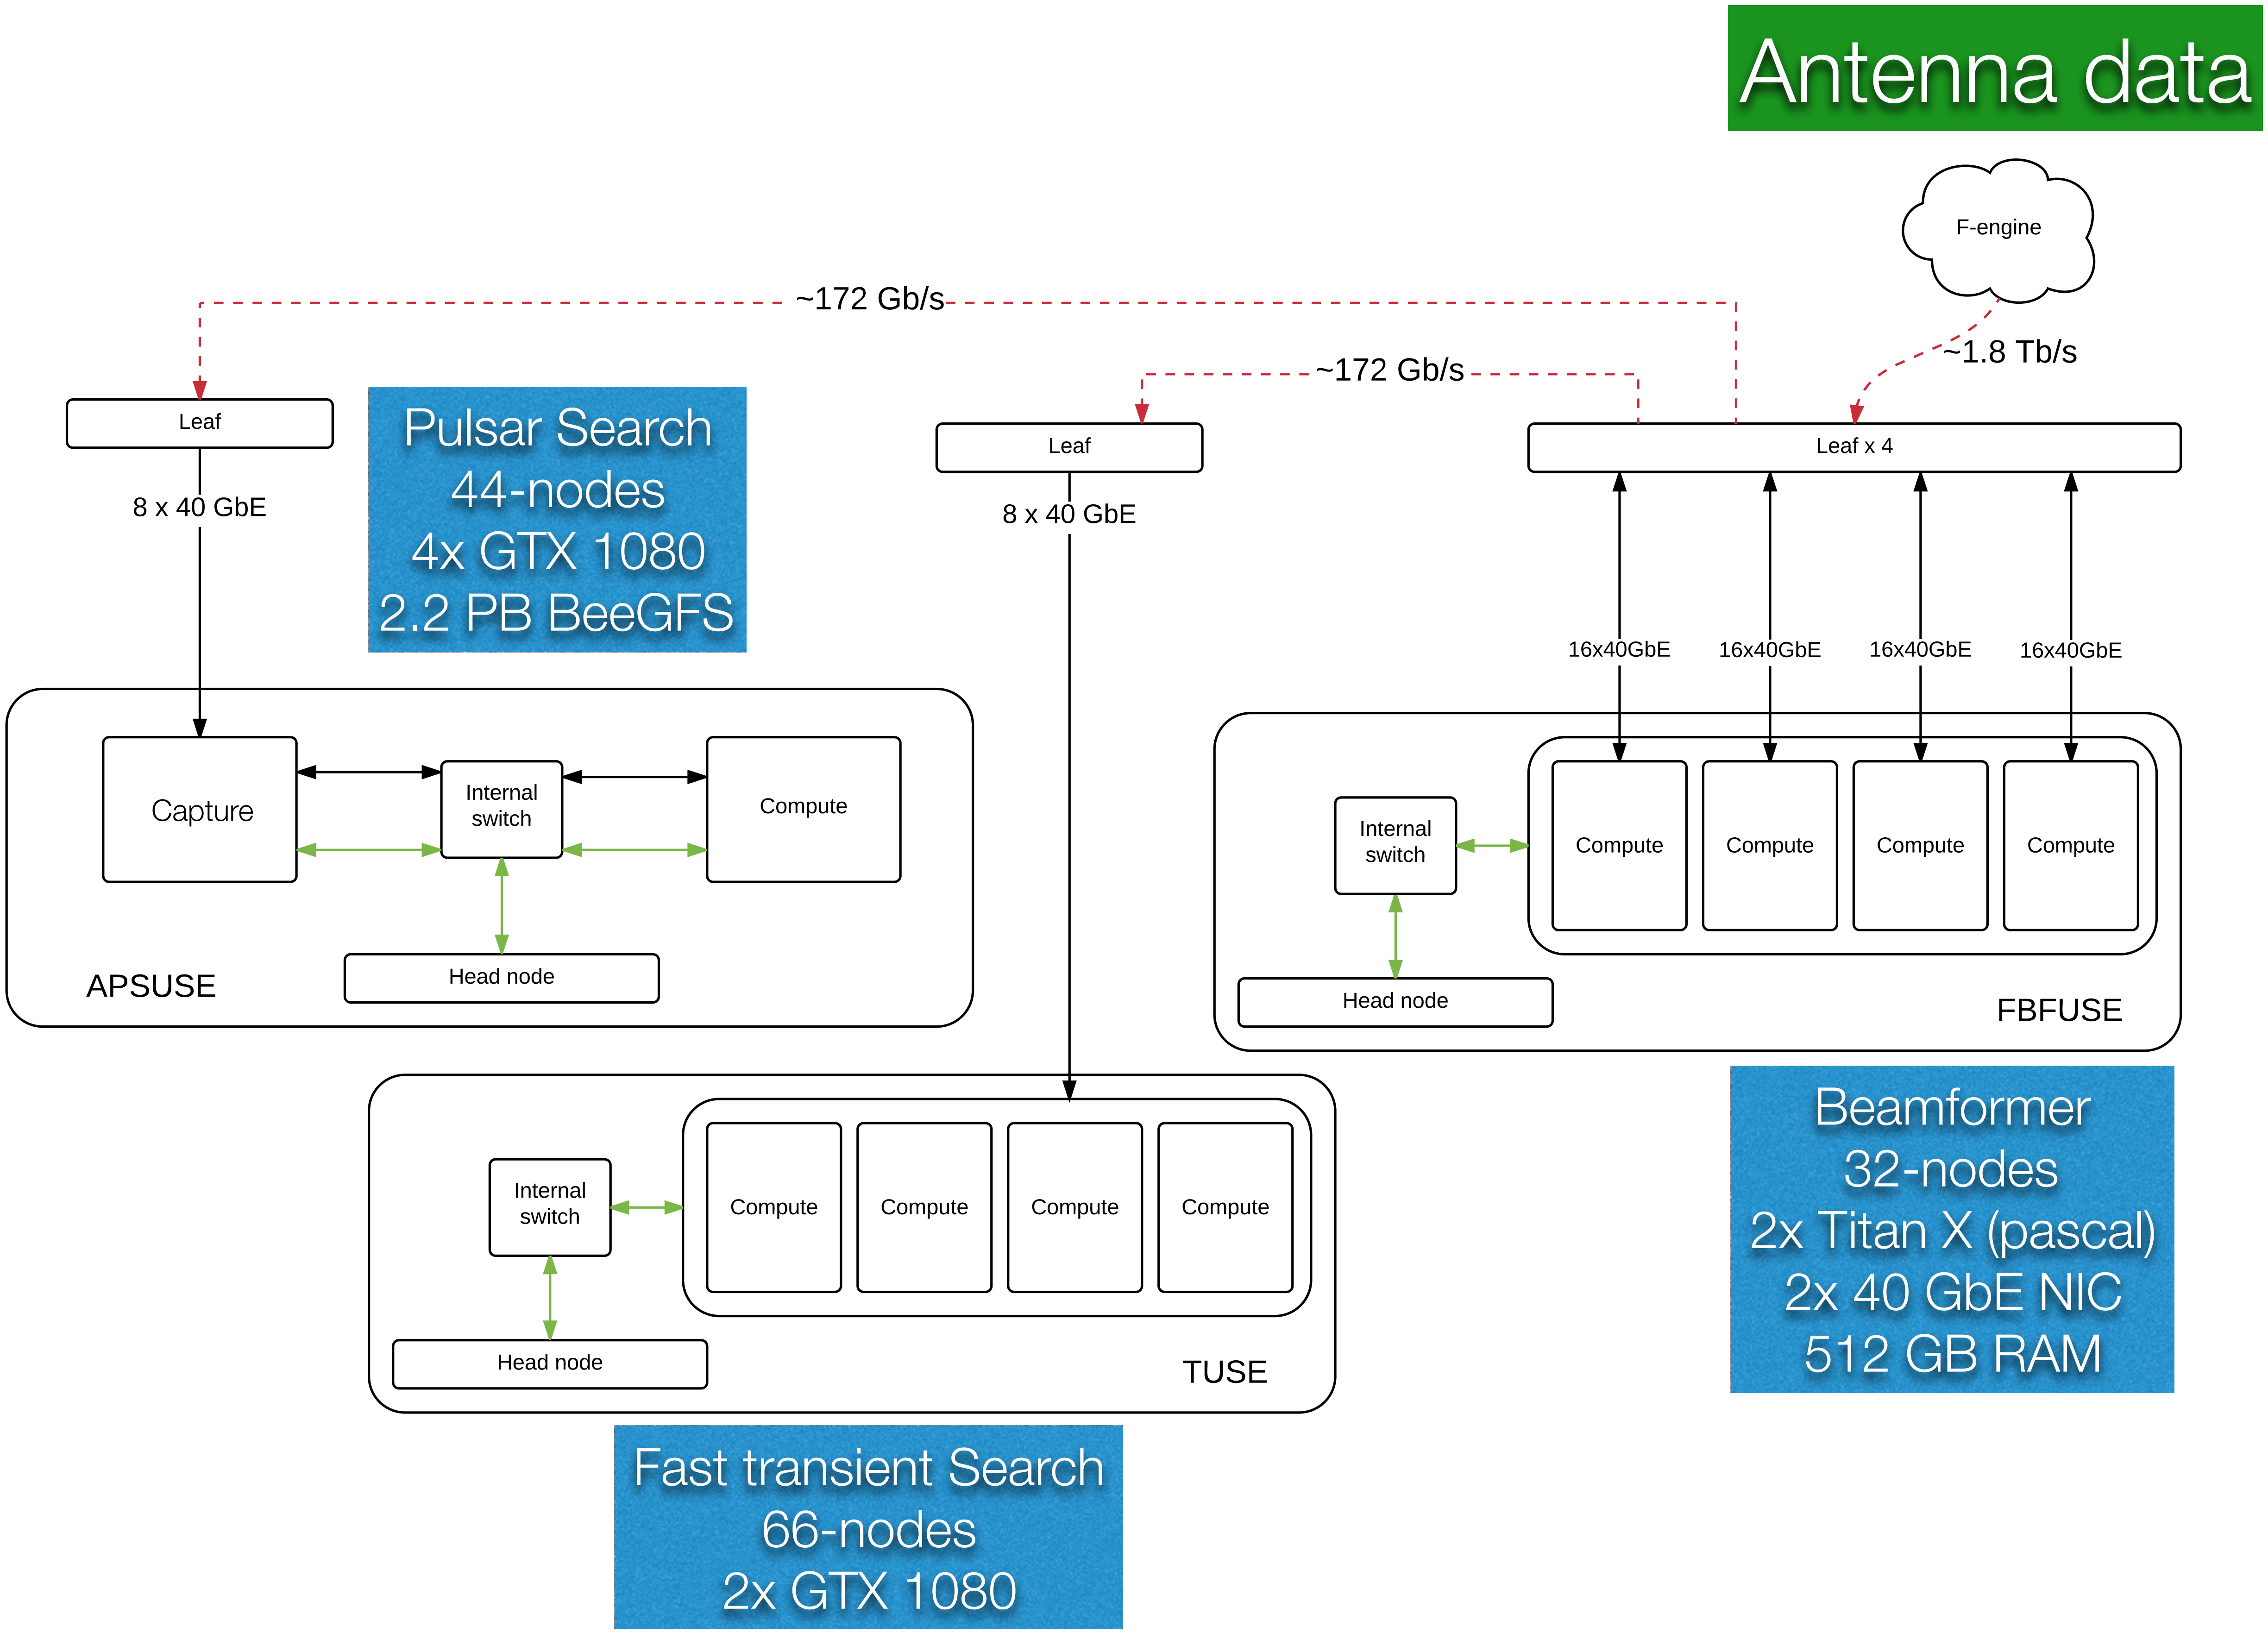
Antenna data
F-engine
~1.8 Tb/s
~172 Gb/s
~172 Gb/s
Leaf
Leaf
Leaf x 4
8 x 40 GbE
8 x 40 GbE
16x40GbE
16x40GbE
16x40GbE
16x40GbE
Pulsar Search
44-nodes
4x GTX 1080
2.2 PB BeeGFS
Capture
Internal switch
Compute
Head node
APSUSE
Internal switch
Compute
Compute
Compute
Compute
Head node
FBFUSE
Internal switch
Compute
Compute
Compute
Compute
Head node
TUSE
Beamformer
32-nodes
2x Titan X (pascal)
2x 40 GbE NIC
512 GB RAM
Fast transient Search
66-nodes
2x GTX 1080

DP?A performance

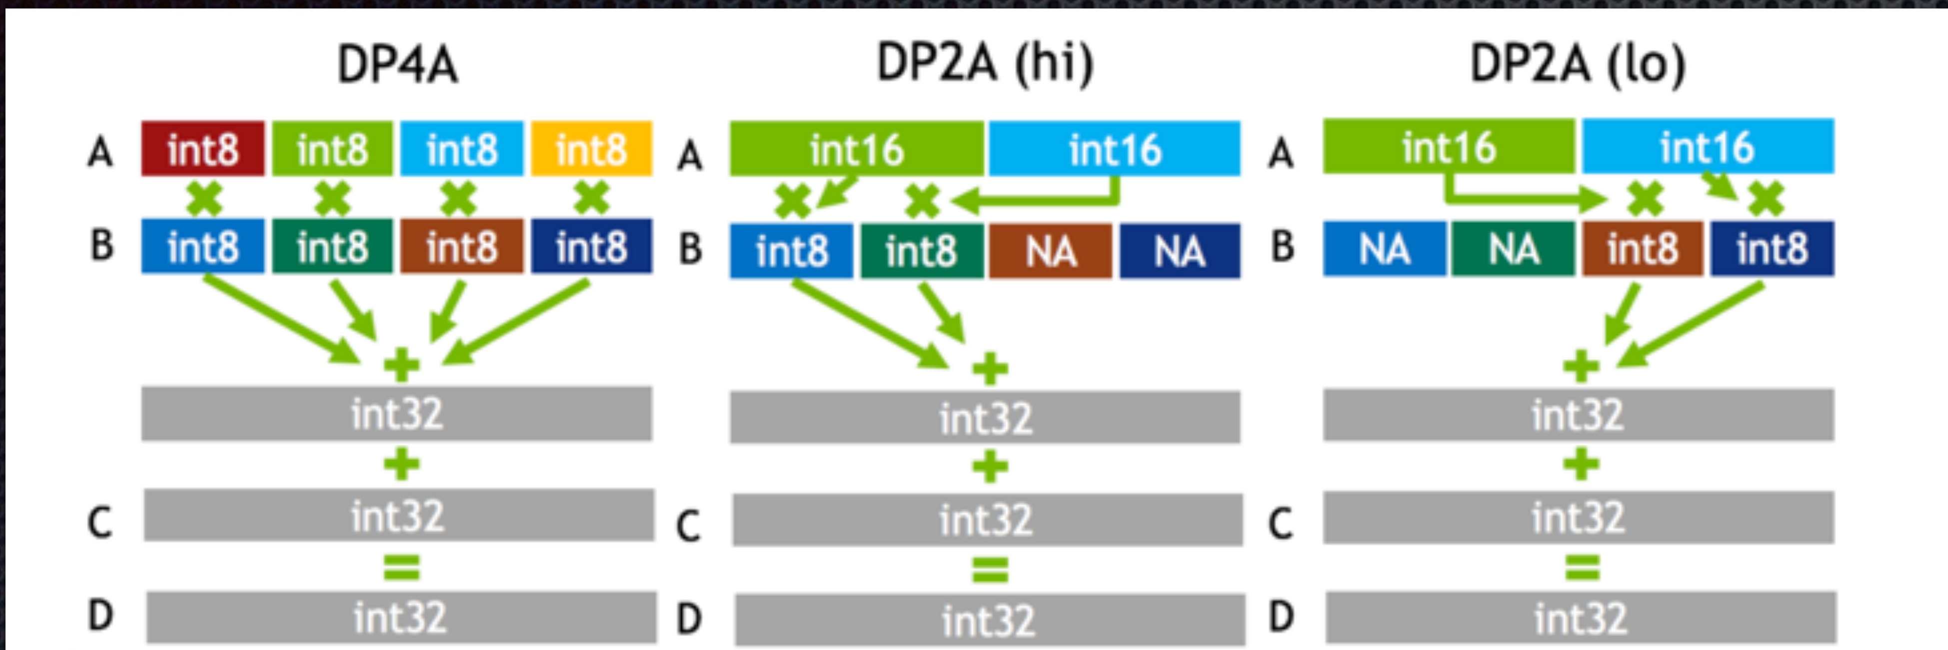


```
__forceinline__ __device__
void dp4a(int &c, const int &a, const int &b) {
  asm("dp4a.s32.s32 %0, %1, %2, %3;" : "+r"(c) : "r"(a), "r"(b), "r"(c));
}
```

- Available on:
  - GP102
  - GP104
  - GP106

- For complex data requires byte-wise transpose
- For beamforming and correlation (complex multiply add) gives phenomenal performance:
  - Without dp4a: **4.1 Tflops**
  - With dp4a: **15.7 Tops**

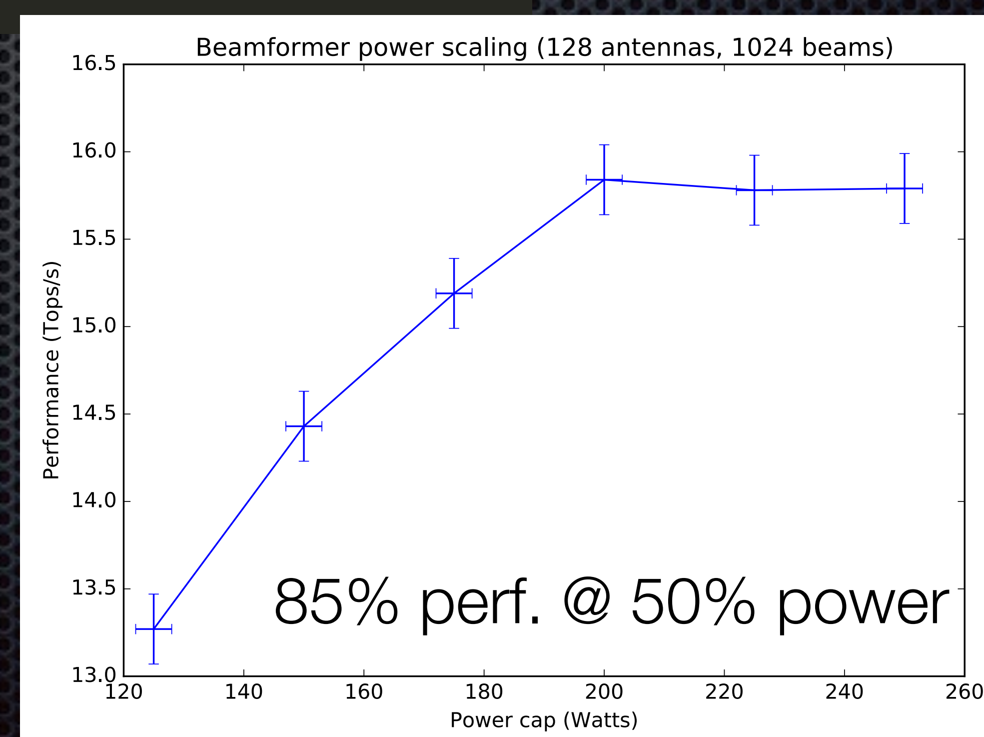

S-band system deployment

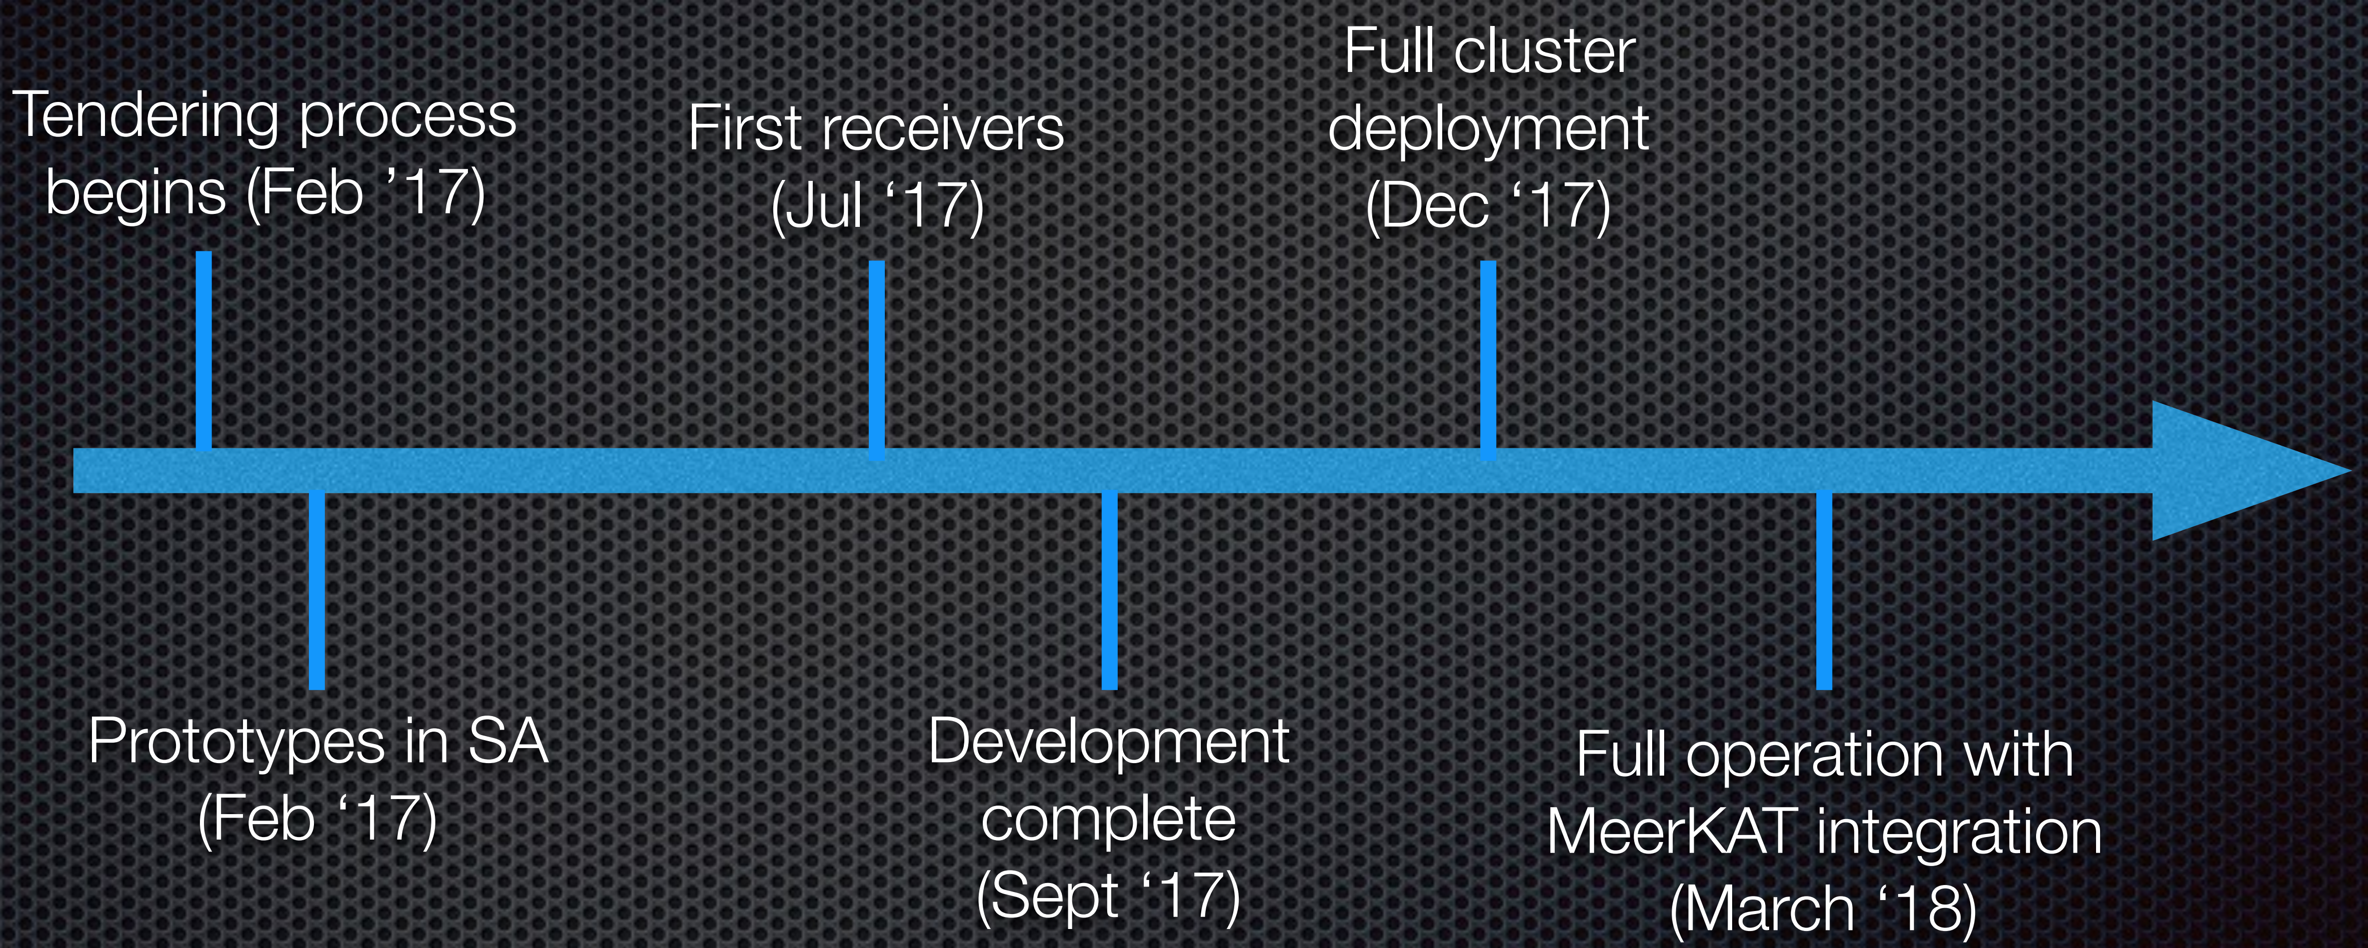

Summary

- Many big data challenges for SKA and its precursors:
  - High processing loads
  - High data throughput
  - Complex scheduling
  - Enormous storage requirements
- Multiple projects are currently ongoing trying to de-risk these issues
- Lots of fantastic science expected
- Great opportunity to hunt for the unknown unknowns!