From Terabytes to Petabytes

Matthew Bailes

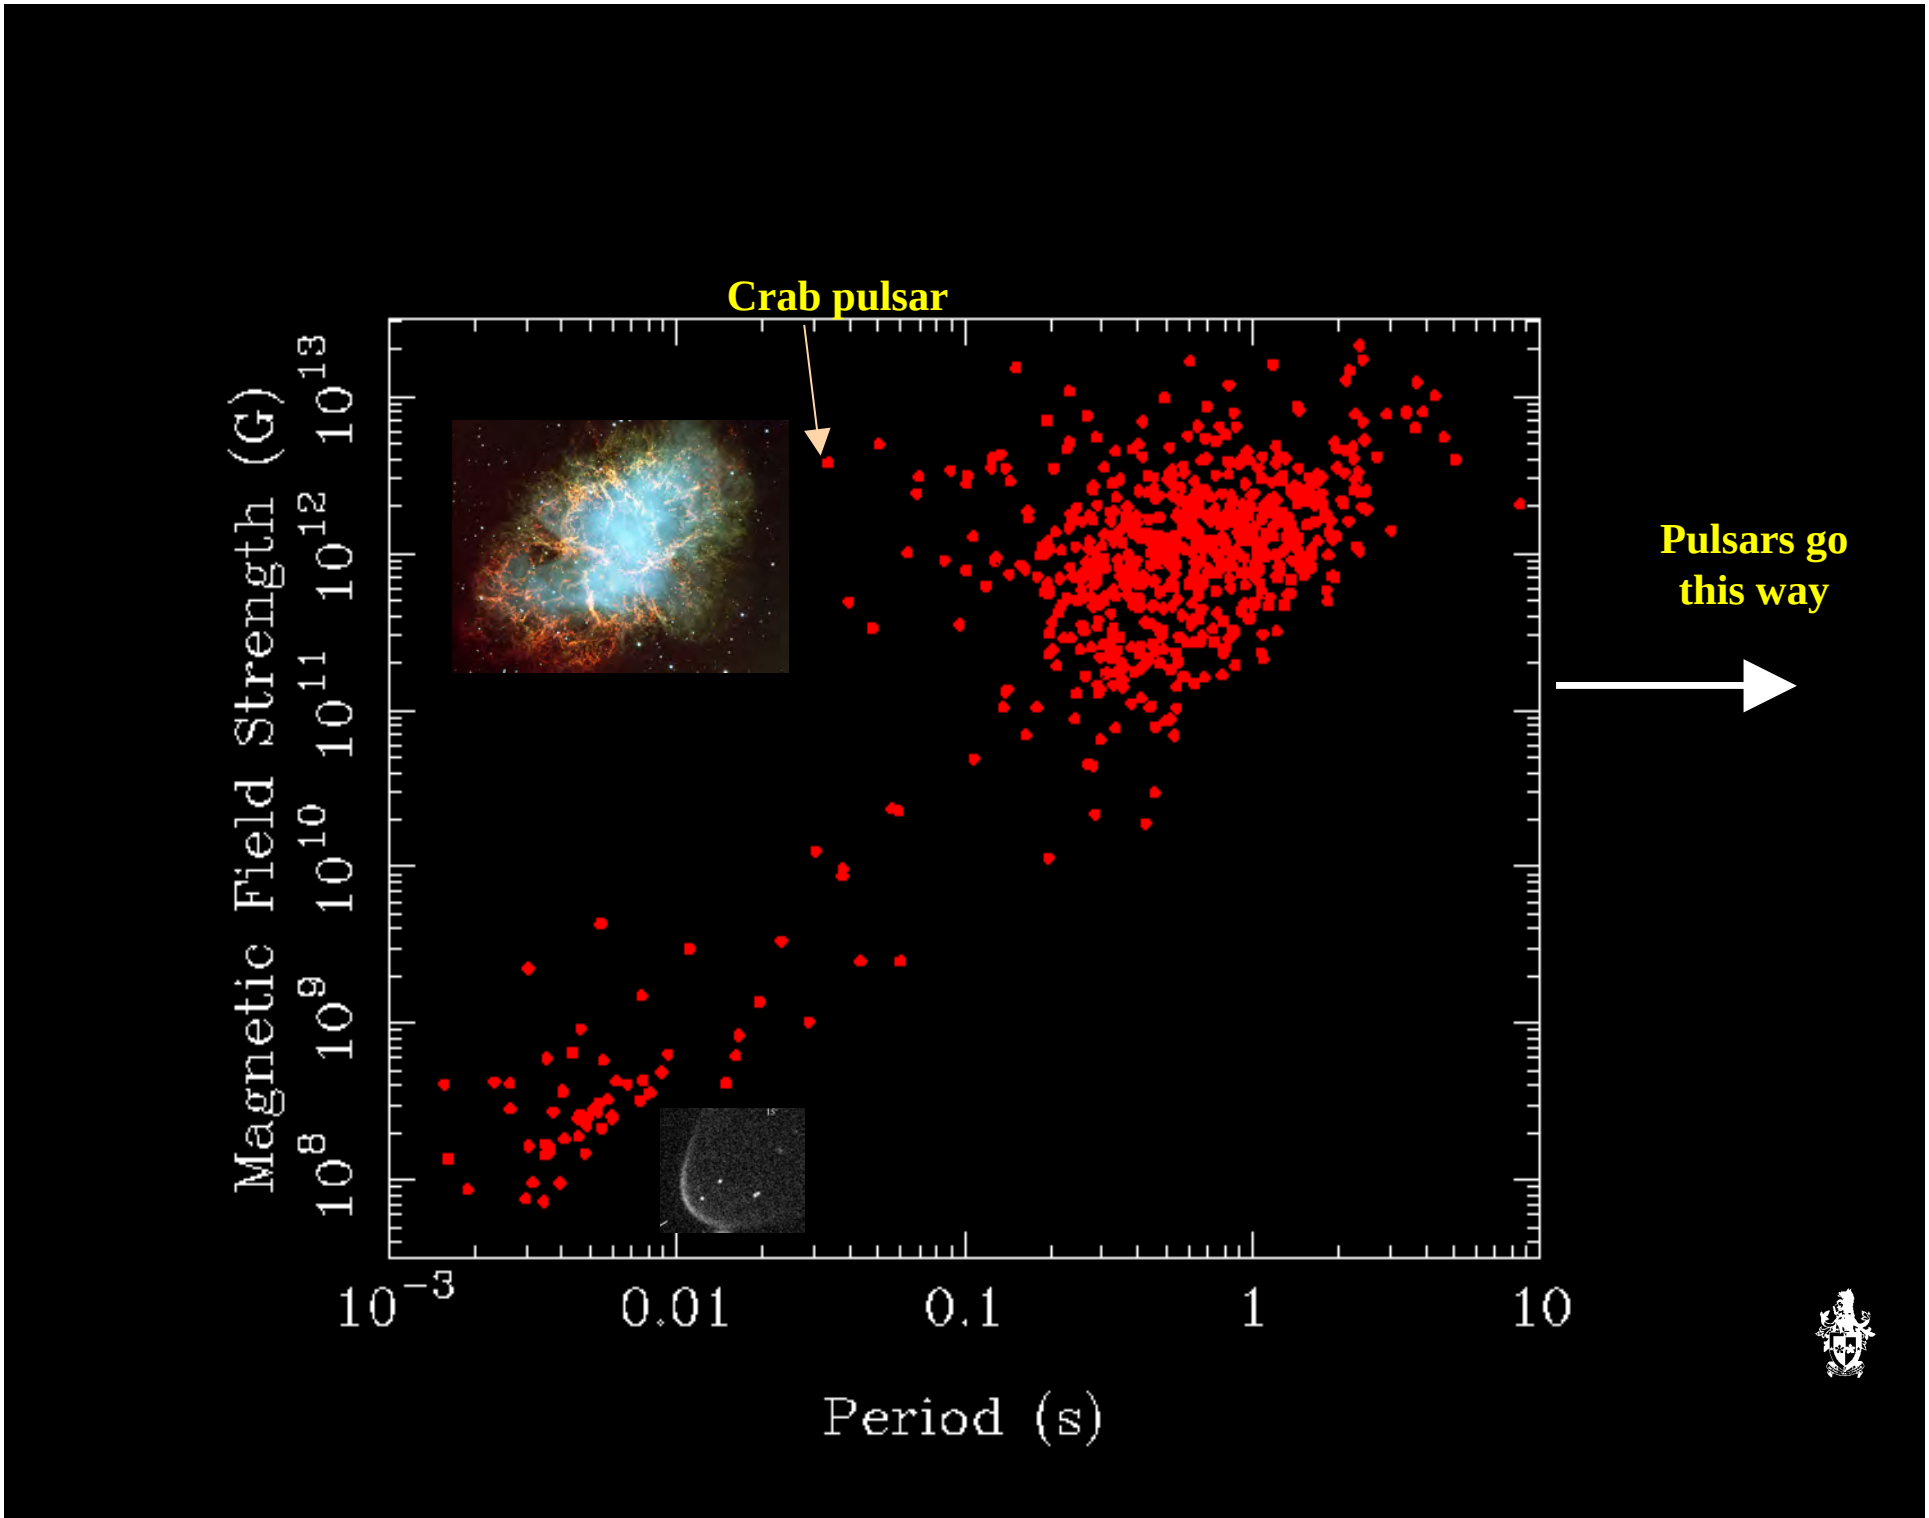
Crab pulsar
Pulsars go this way
Magnetic Field Strength (G)
10^{8} 10^{9} 10^{10} 10^{11} 10^{12} 10^{13}
10^{-3} 0.01 0.1 1 10
Period (s)

One-bit Parkes Multibeam Pulsar Surveys

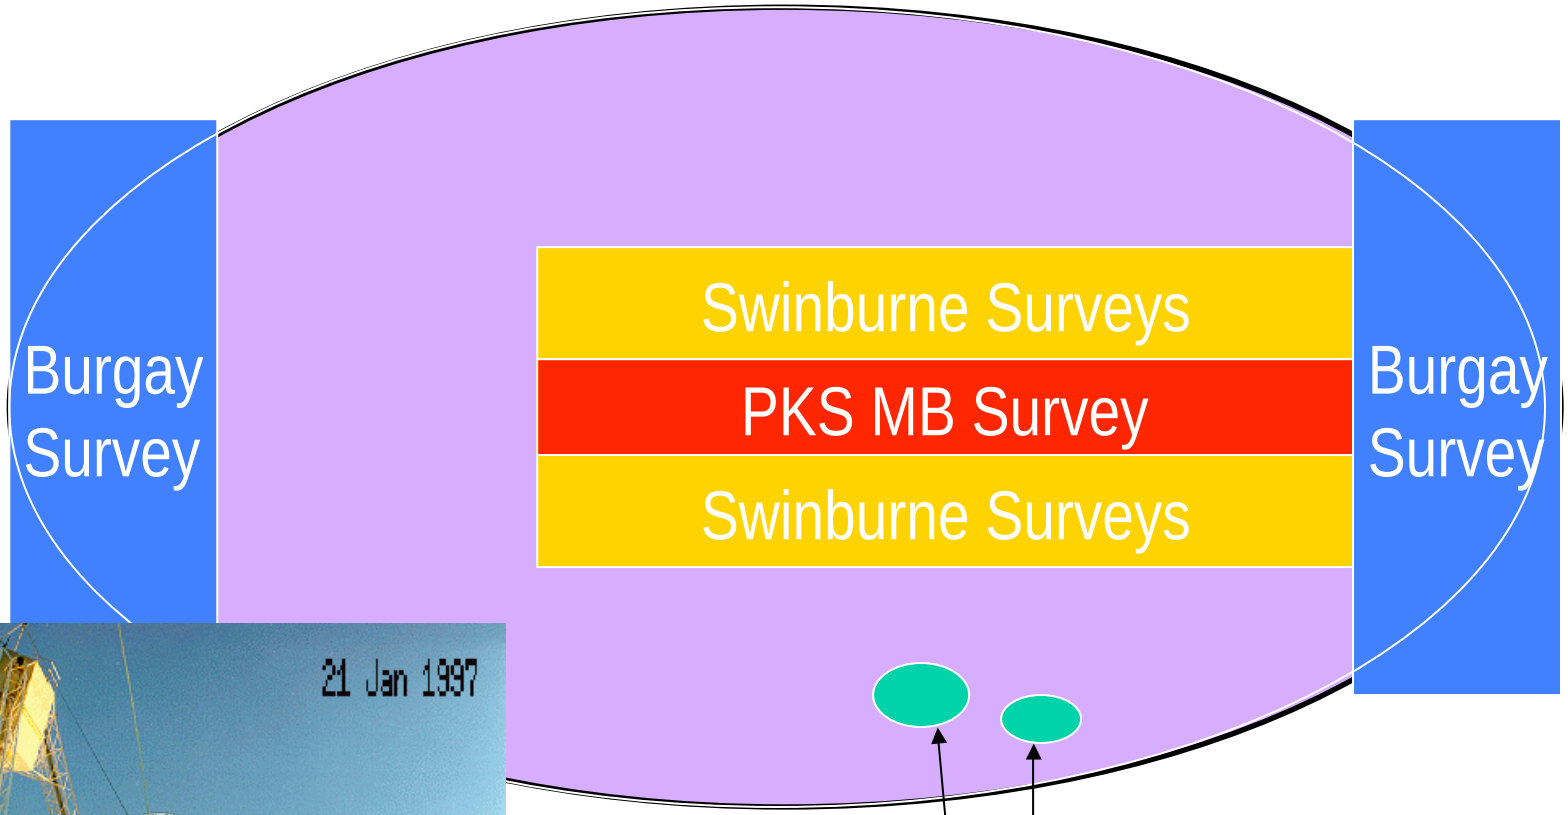

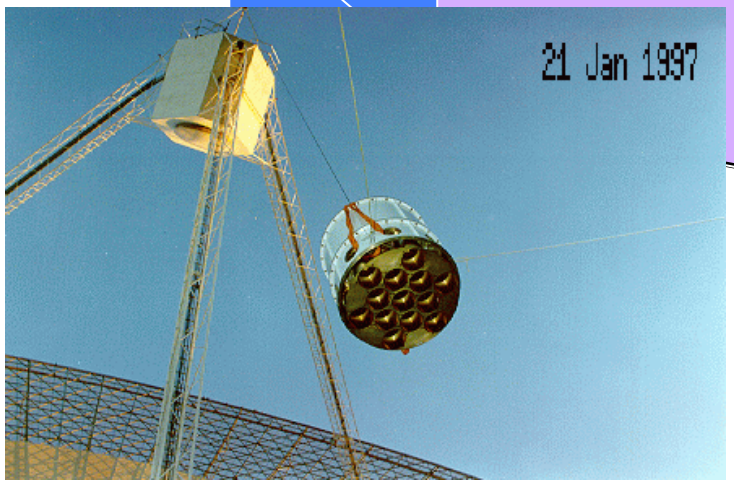

Old Data Rates

- 13 x 96 channels x 1 bit every 250 us = 624 kB/s.
- 35 min pointings
- < 5 TB

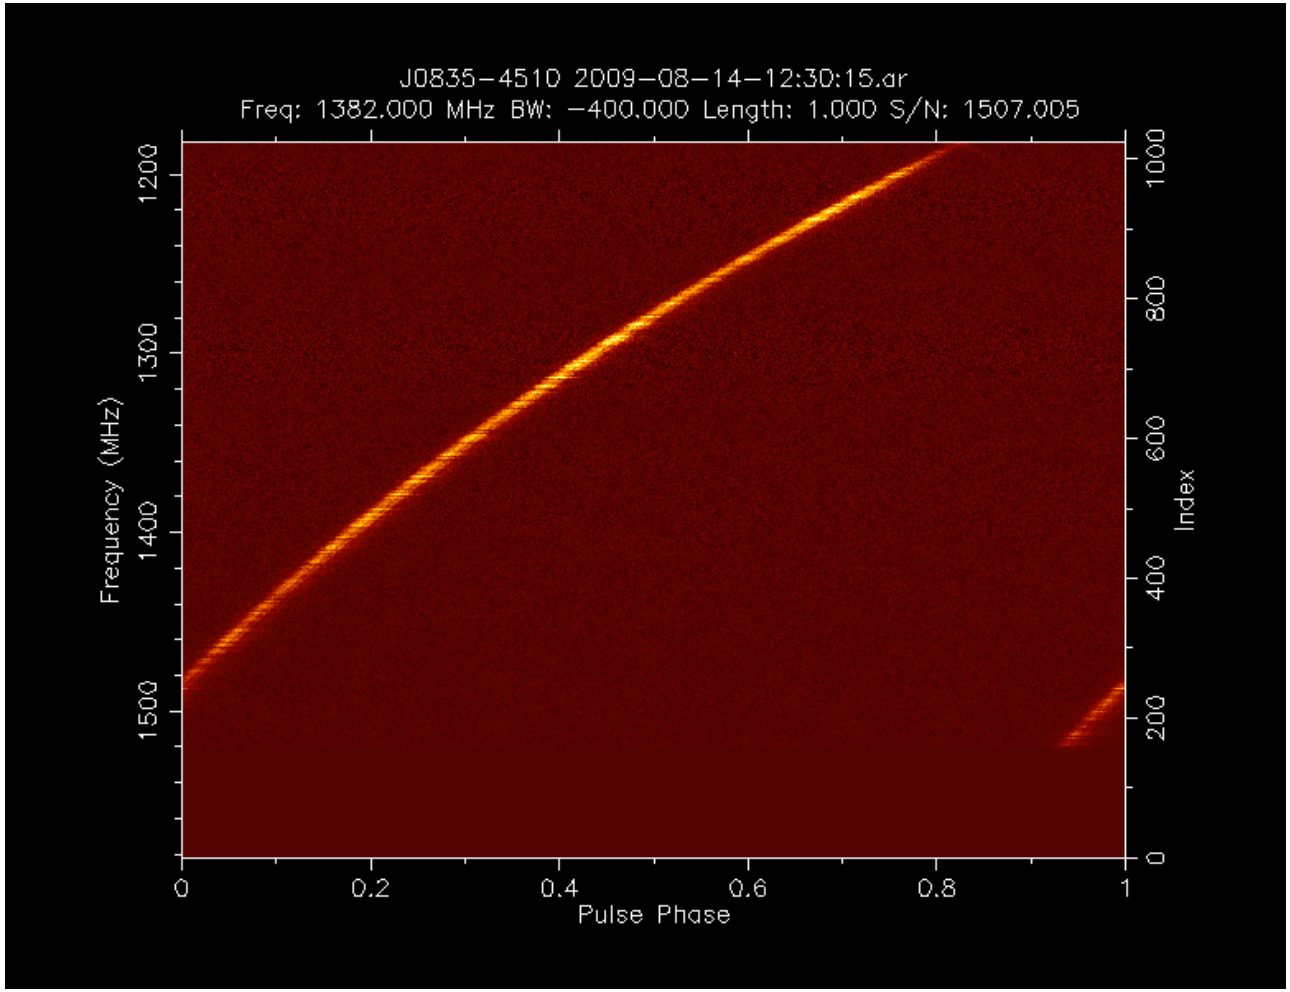

SWINBURNE UNIVERSITY OF TECHNOLOGY

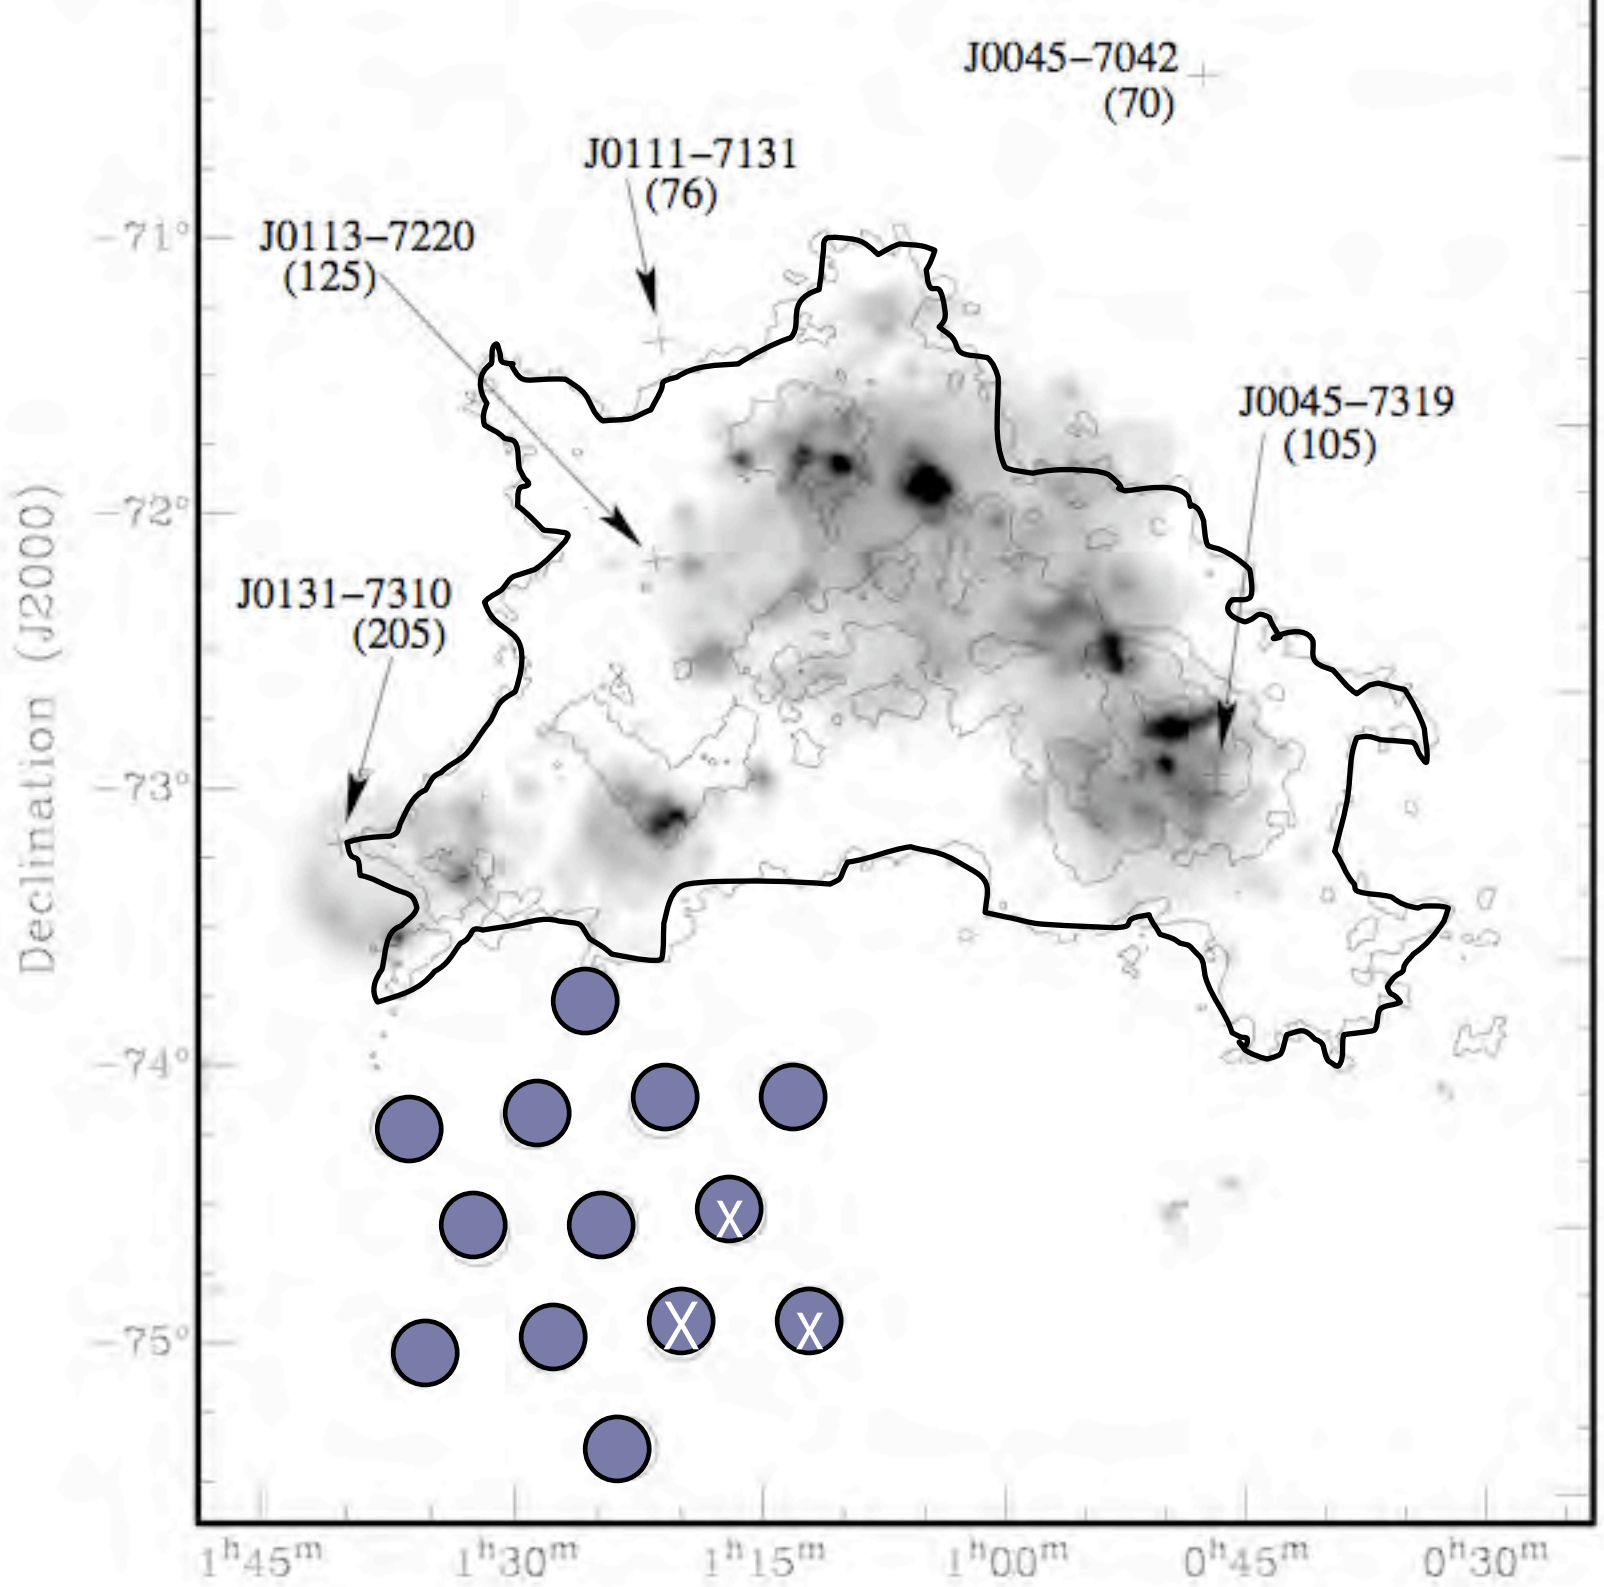
J0045−7042
(70)
J0111−7131
(76)
J0113−7220
(125)
J0045−7319
(105)
J0131−7310
(205)
−71°
−72°
−73°
−74°
−75°
Declination (J2000)
1h45m
1h30m
1h15m
1h00m
0h45m
0h30m

LORIMER BURST (2001 data, published in Lorimer et al. 2007)

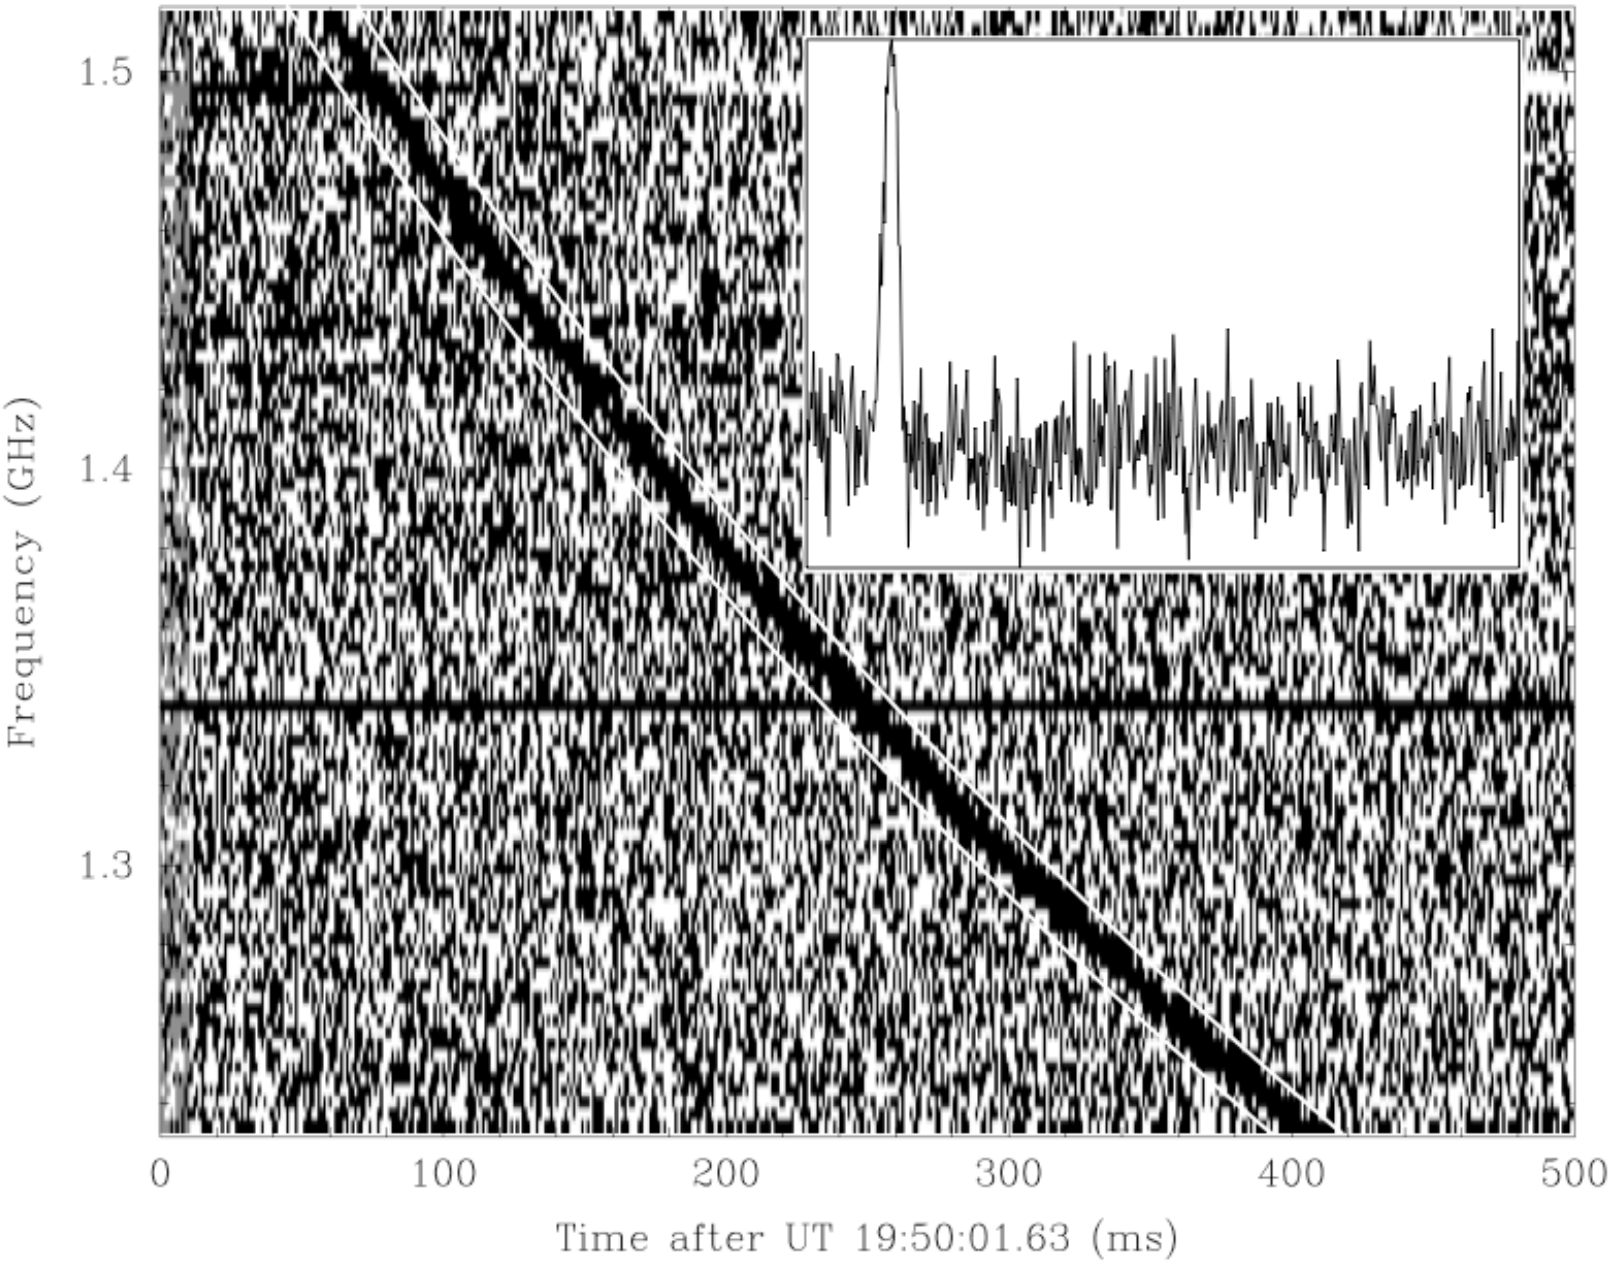

Old 1-bit Hardware

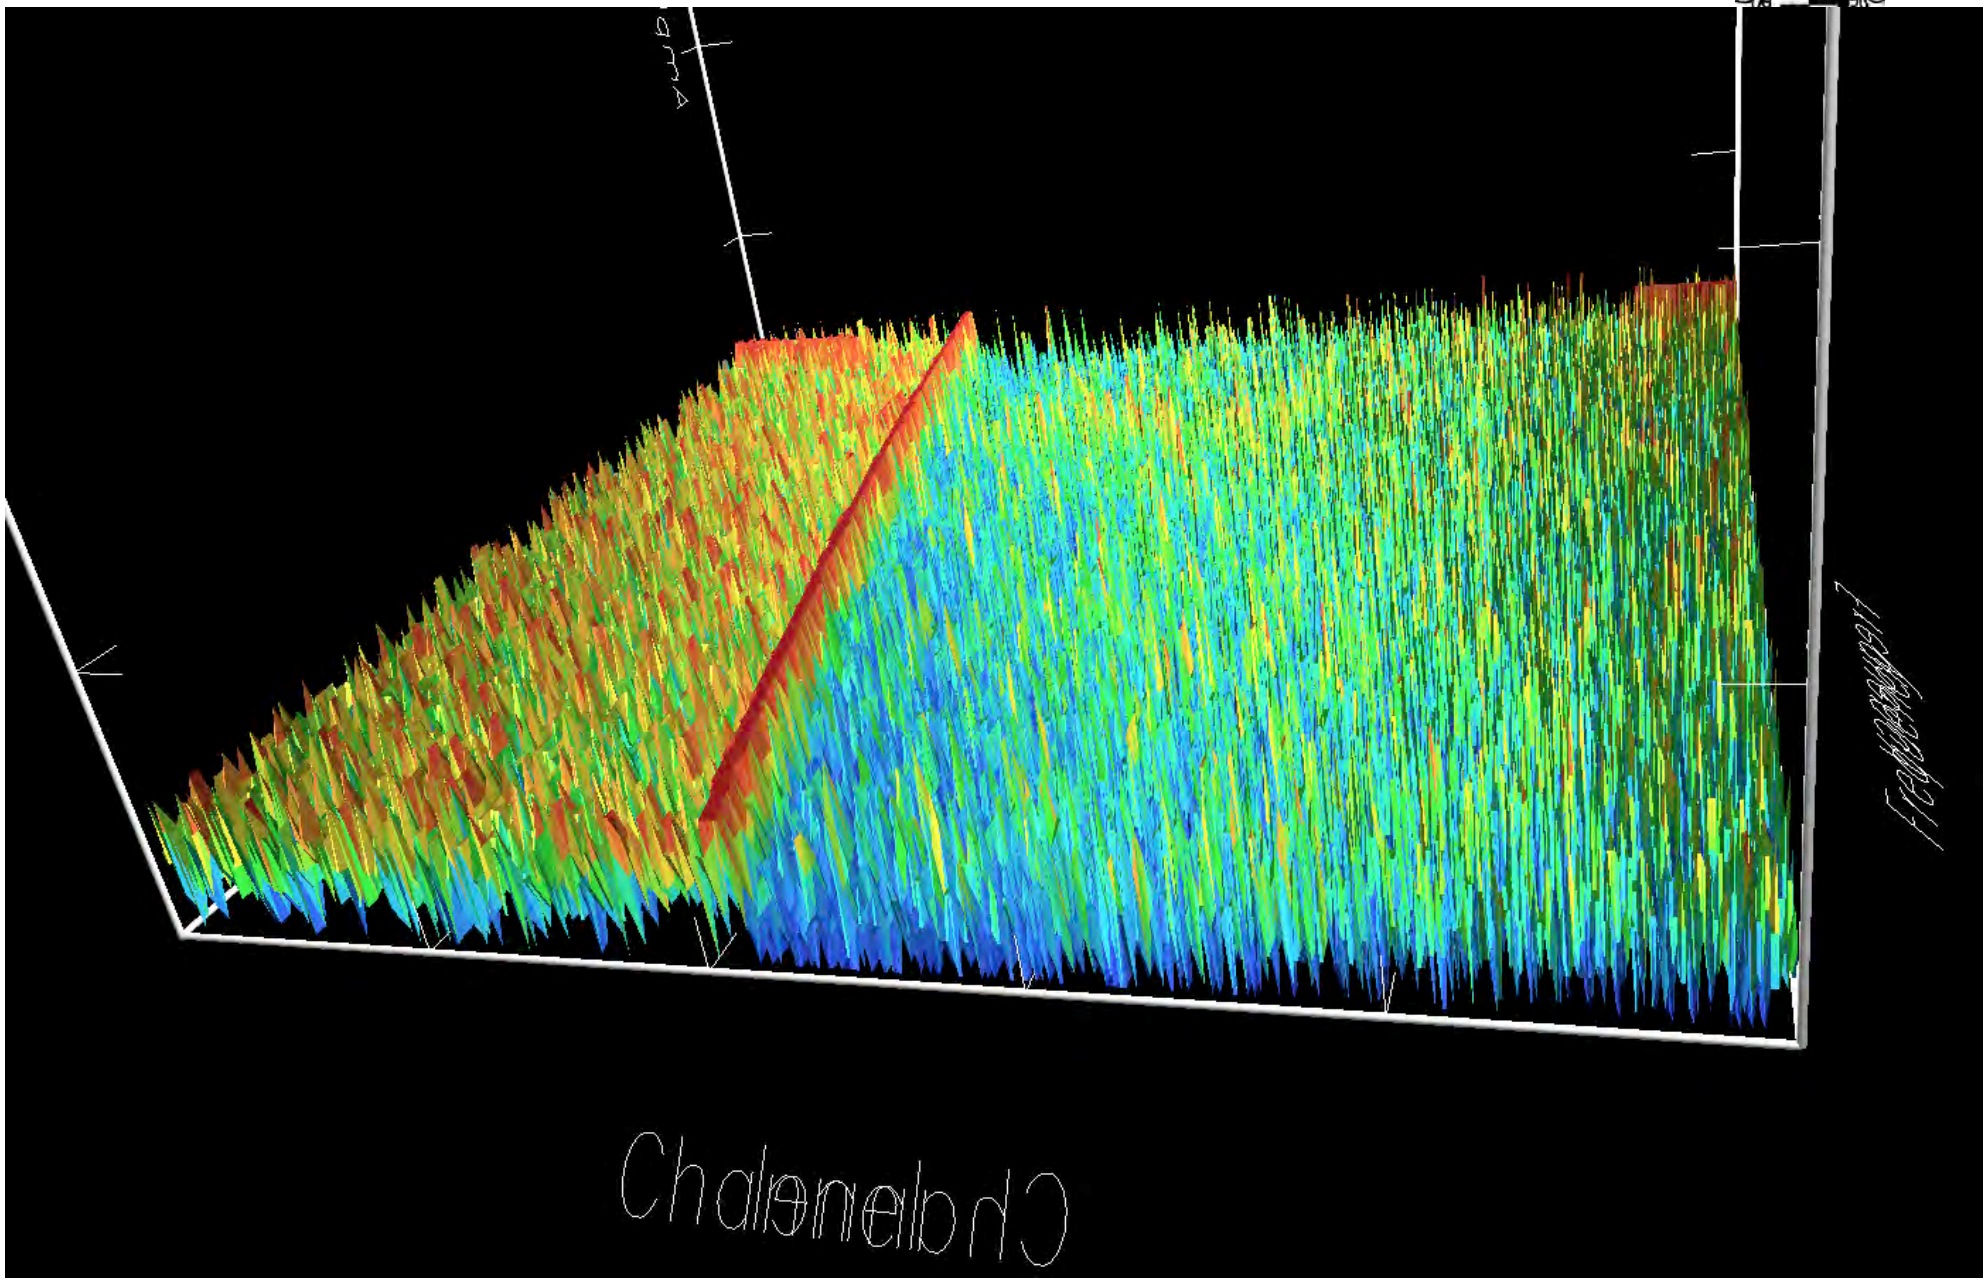


CHEAP Digital Signal Processing
ROACH/iBOB CASPER Tools
Staveley-Smith & Bailes
Danny Price
Andrew Jameson
“HIPSR”

From Terabytes to Petabytes

The High Time Resolution Universe Pulsar Survey – I. System configuration and initial discoveries

M. J. Keith,[1★] A. Jameson,[2] W. van Straten,[2] M. Bailes,[2] S. Johnston,[1] M. Kramer,[3,4] A. Possenti,[5] S. D. Bates,[4] N. D. R. Bhat,[2] M. Burgay,[5] S. Burke-Spolaor,[1,2] N. D'Amico,[5] L. Levin,[1,2] Peter L. McMahon,[6,7] S. Milia[5,8] and B. W. Stappers[4]

[1]*Australia Telescope National Facility, CSIRO, PO Box 76, Epping, NSW 1710, Australia*
[2]*Swinburne University of Technology, Centre for Astrophysics and Supercomputing, Mail H39, PO Box 218, VIC 3122, Australia*
[3]*Max Planck Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany*
[4]*University of Manchester, Jodrell Bank Centre for Astrophysics, Alan Turing Building, Manchester M13 9PL*
[5]*INAF – Osservatorio Astronomico di Cagliari, località Poggio dei Pini, strada 54, I-09012 Capoterra, Italy*
[6]*Berkeley Wireless Research Center, University of California, Berkeley, CA 94704, USA*
[7]*Department of Electrical Engineering, University of Cape Town, Rondebosch, 7701, South Africa*
[8]*Dipartimento di Fisica, Università degli Studi di Cagliari, Cittadella Universitaria, 09042 Monserrato (CA), Italy*

Accepted 2010 July 7. Received 2010 June 28; in original form 2010 May 11

SWINBURNE UNIVERSITY OF TECHNOLOGY

HITRUN Pulsar Surveys (Keith et al. 2010)

- Digital Versions of the Parkes MB Surveys
- 13 x 1024 x 64 us
  - Medium-Latitude Survey (540s)
  - Ultra-deep Survey (4000s)
  - All-sky Survey (260s)
- Better for high-DM MSPs
- Better for Lorimer bursts
- Better for RRATs
- Deeper than old surveys

1 Petabyte of Data!

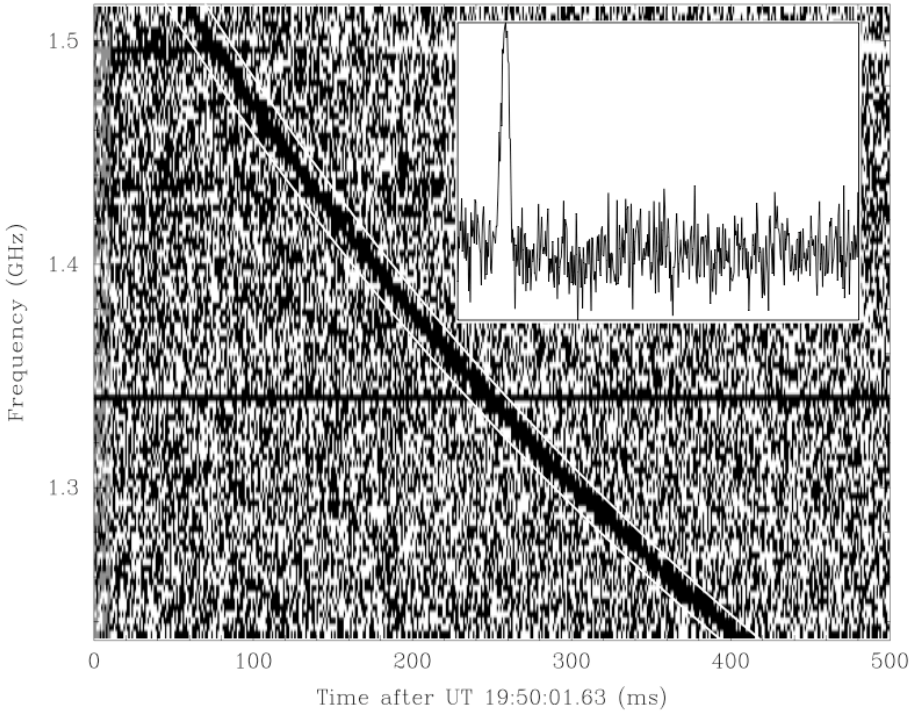

SWINBURNE UNIVERSITY OF TECHNOLOGY

Data Volumes

SWINBURNE UNIVERSITY OF TECHNOLOGY

SWINBURNE PULSAR PORTAL

HTRU Survey
- Survey Summary

Search Database
- Range Search
- Cone Search

Missing Obs
- Medlat
- Hilat
- All Survey
- All Obs

Processing
- Summary
- View Jobs **

Account
- Account Home
- User Settings
- Logout

HTRU Survey Summary

Total Observations	Ultradeep	Med Lat	High Lat	Total
Number of Observations	1226	7756	35474	43230
Number of Beams	15922	100672	460356	561028
Size [TB]	258.23	212.28	505.51	976.02

On Disk	Low Lat	Med Lat	High Lat	Total
Number of Observations	413	7706	23318	31024
Number of Beams	4570	99941	302629	402570
Size [TB]	74.13	210.73	333.70	544.43
% Total	28.71	99.27	66.01	75.85

Issues:

- To read from tape:
  - 50 MB/s
  - T_read = 1e15/50e6 = 231 days!
- Manually, every 4 hours, 3 tapes/day
  - 462 days
- Weekends / (5/7)
  - 646 days!
- Use Robots
  - Jam, expensive, need cleaning.

SWINBURNE UNIVERSITY OF TECHNOLOGY

Disk Space

- High performance RAID ~ 1M$ / PB.
  - Sustains 700 MB/s on read
- T_read = 16.5 days!
- Cannot backup such a disk cheaply
- Mirroring (1M)
- MAID (0.4M)
- Tape system (500K?)

SWINBURNE UNIVERSITY OF TECHNOLOGY

Job Processing

- Originally: 4 h/job, 95,000 jobs for 25% processing
- One machine:
  - T = 43 machine years (medlat) or 200 years (all)
- Buy 150 machines! (green supercomputer)

SWINBURNE UNIVERSITY OF TECHNOLOGY

Batch Queue

- Doesn't like 95,000 jobs
- Read speed only 30 MB/s (nfs)
- IO issues when multiple accesses
- Scenario:

Sharing with cosmologist running on 100 nodes, job ends, 100 HTRU jobs start at once, all hit disk, disk breaks.

SWINBURNE UNIVERSITY OF TECHNOLOGY

Solutions

- Never submit more than one job/read time.
- Touch an “access” file, job sleeps until disk isn’t busy.
- Parallelise the storage solution.
- New supercomputer. “Gstar”.
  - 700 MB/s IO
  - 2 PB disk
  - >20x processing power (GPUs)

SWINBURNE UNIVERSITY OF TECHNOLOGY

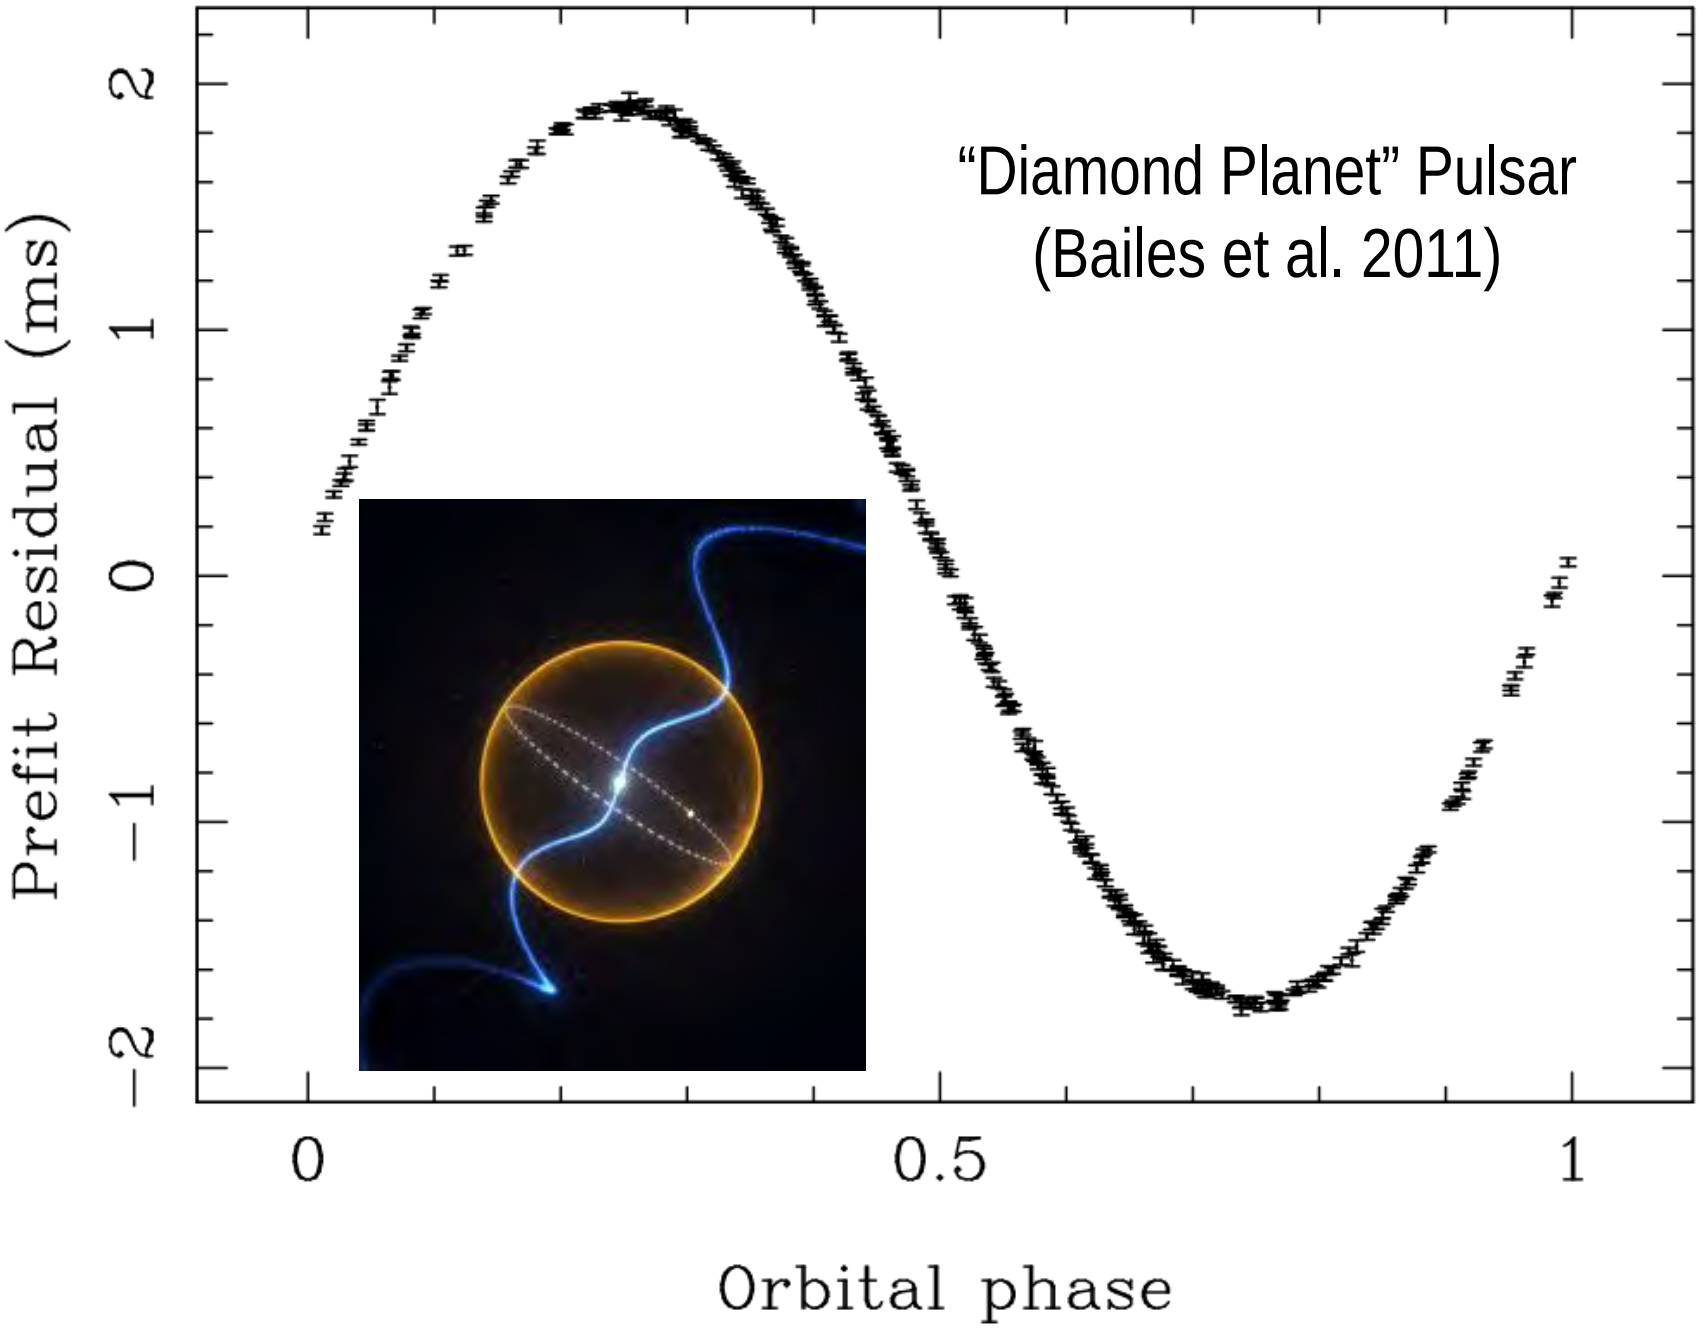
“Diamond Planet” Pulsar
(Bailes et al. 2011)
Prefit Residual (ms)
2
1
0
−1
−2
0
0.5
1
Orbital phase

Fastest MSP in the disk

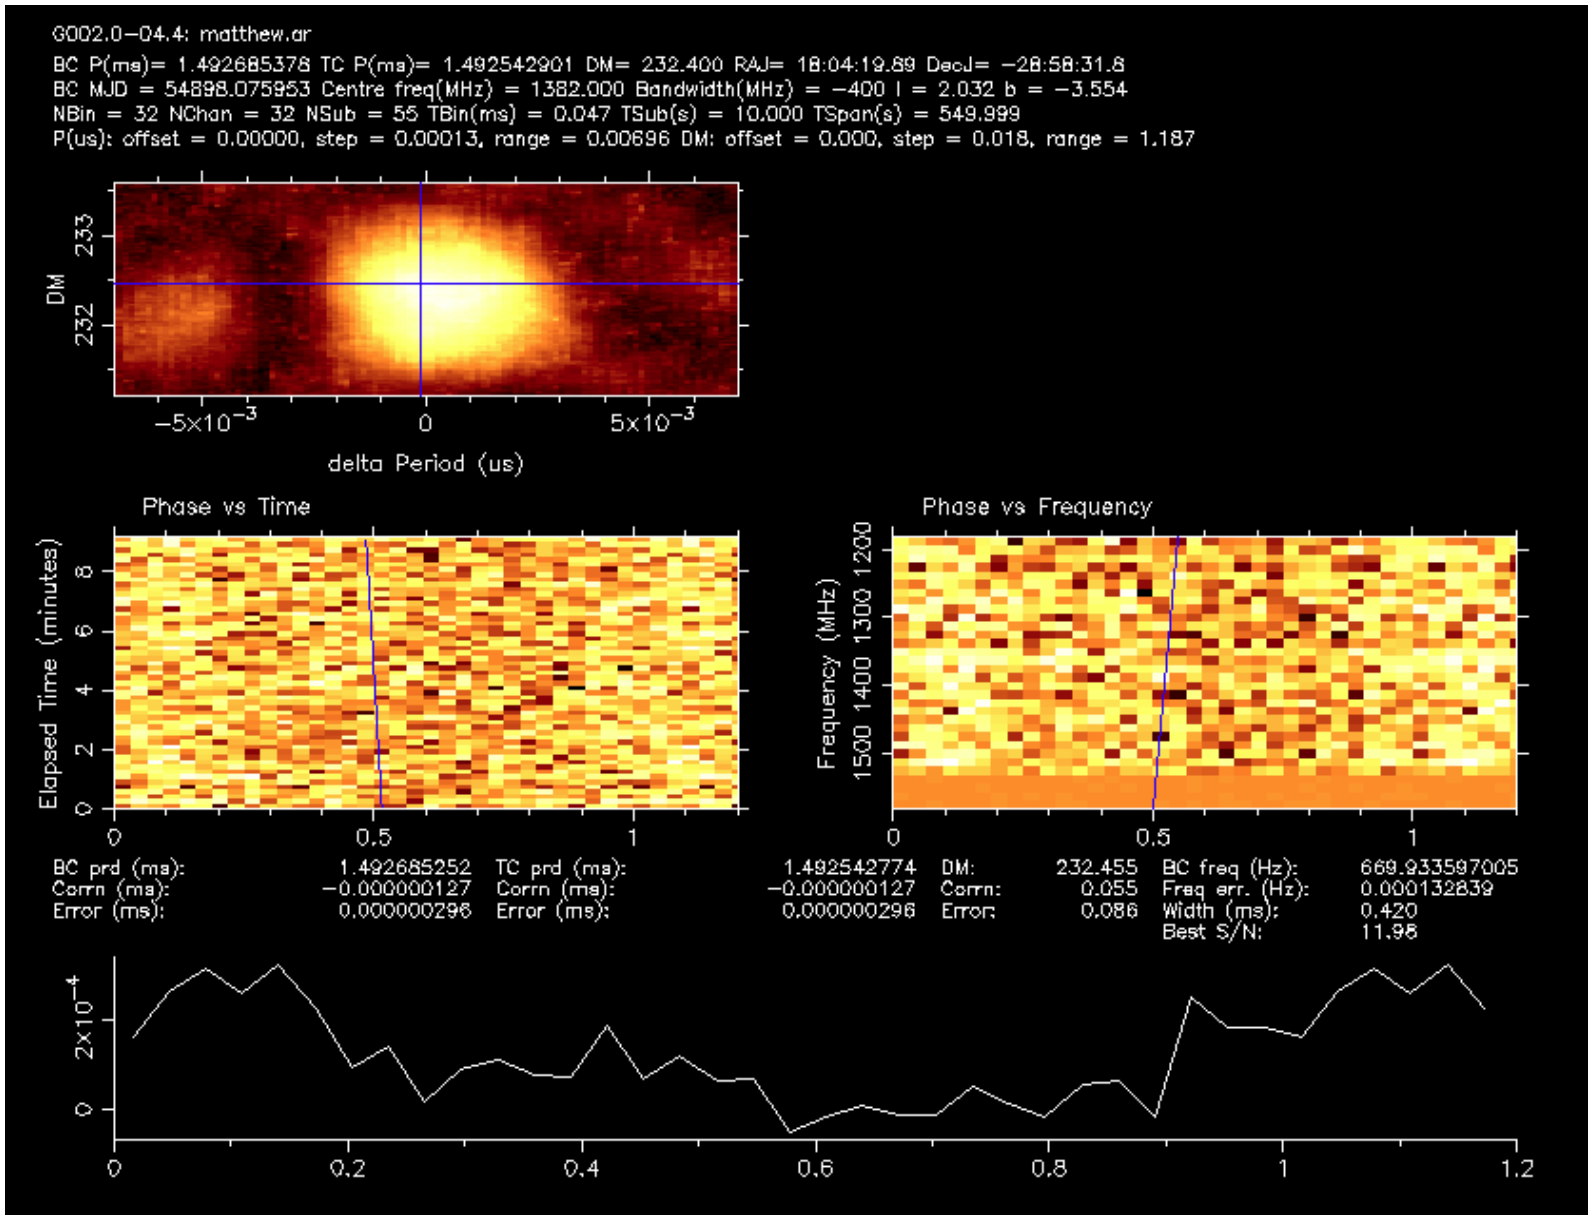

Accelerated MSP

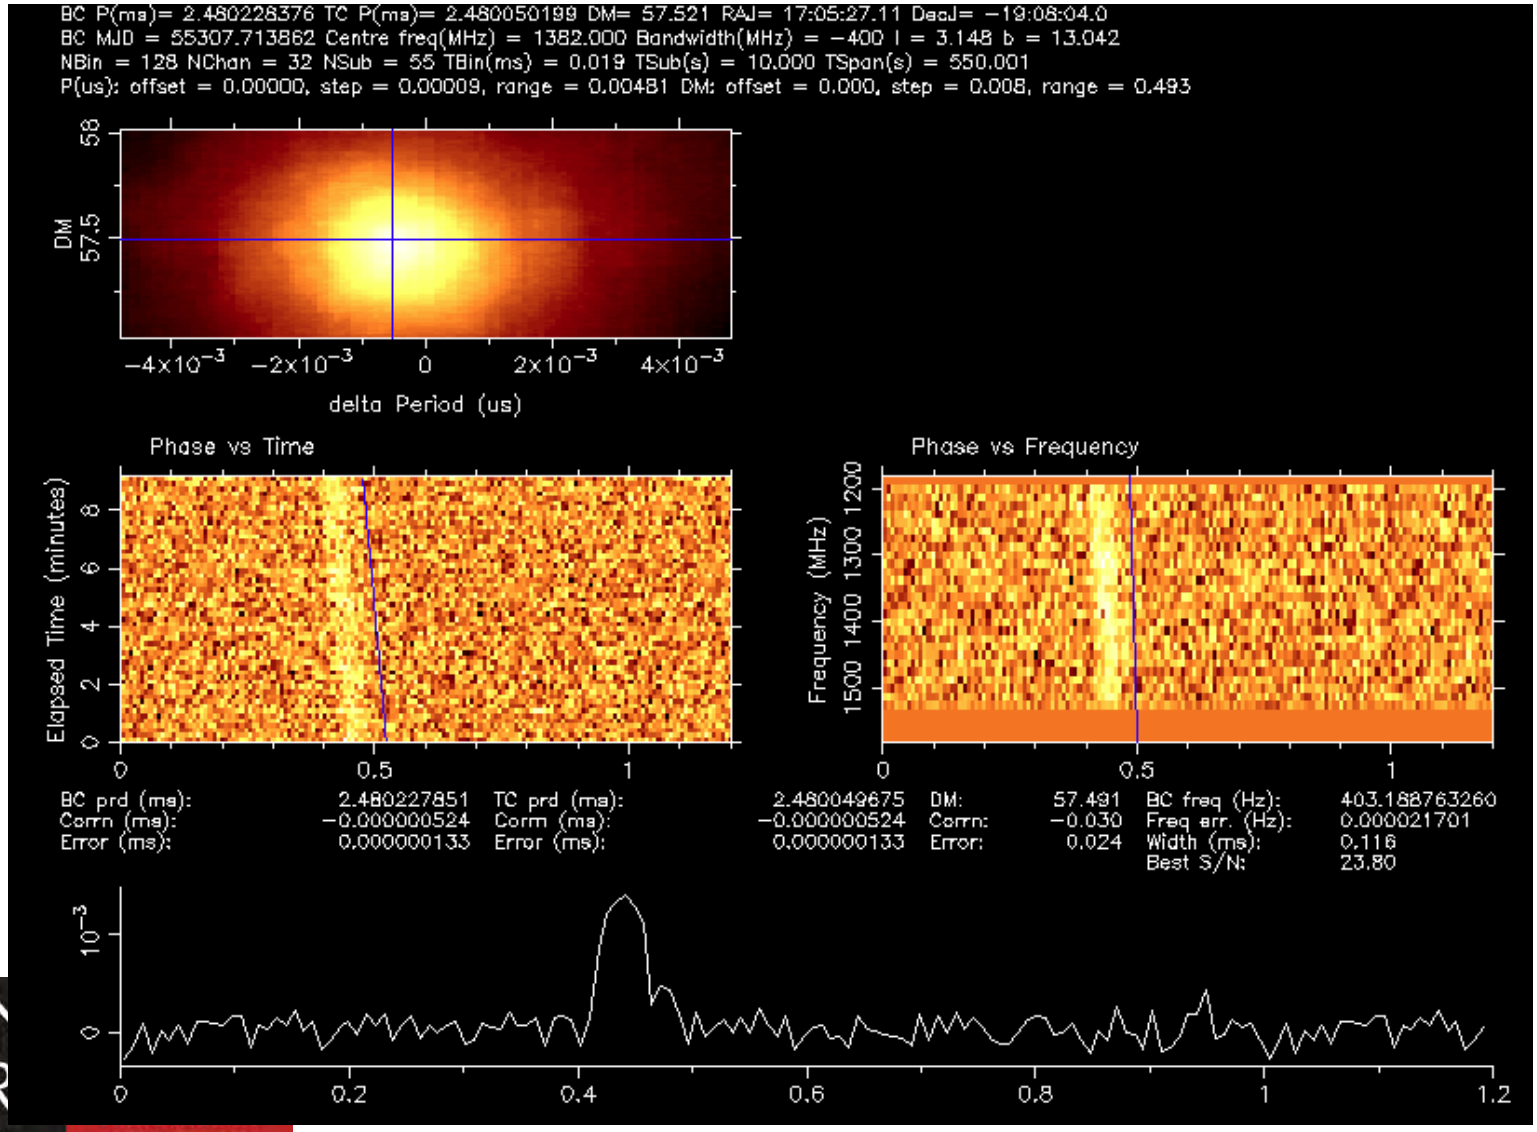

Swinburne Pulsar Portal

SWINBURNE UNIVERSITY OF TECHNOLOGY

SWINBURNE PULSAR PORTAL

HTRU Survey
- Survey Summary

Search Database
- Range Search
- Cone Search

Missing Obs
- Medlat
- Hilat
- All Survey
- All Obs

Processing
- Summary
- View Jobs **

Account
- Account Home
- User Settings
- Logout

HTRU Survey Summary

Total Observations	Ultradeep	Med Lat	High Lat	Total
Number of Observations	1226	7756	35474	43230
Number of Beams	15922	100672	460356	561028
Size [TB]	258.23	212.28	505.51	976.02

On Disk	Low Lat	Med Lat	High Lat	Total
Number of Observations	413	7706	23318	31024
Number of Beams	4570	99941	302629	402570
Size [TB]	74.13	210.73	333.70	544.43
% Total	28.71	99.27	66.01	75.85

SWINBURNE UNIVERSITY OF TECHNOLOGY

DM >> Lorimer Burst ! Well off galactic plane.

One beam
Only.
Instrumentation
better
Allows detailed
studies

DM~944.38pc/cc!

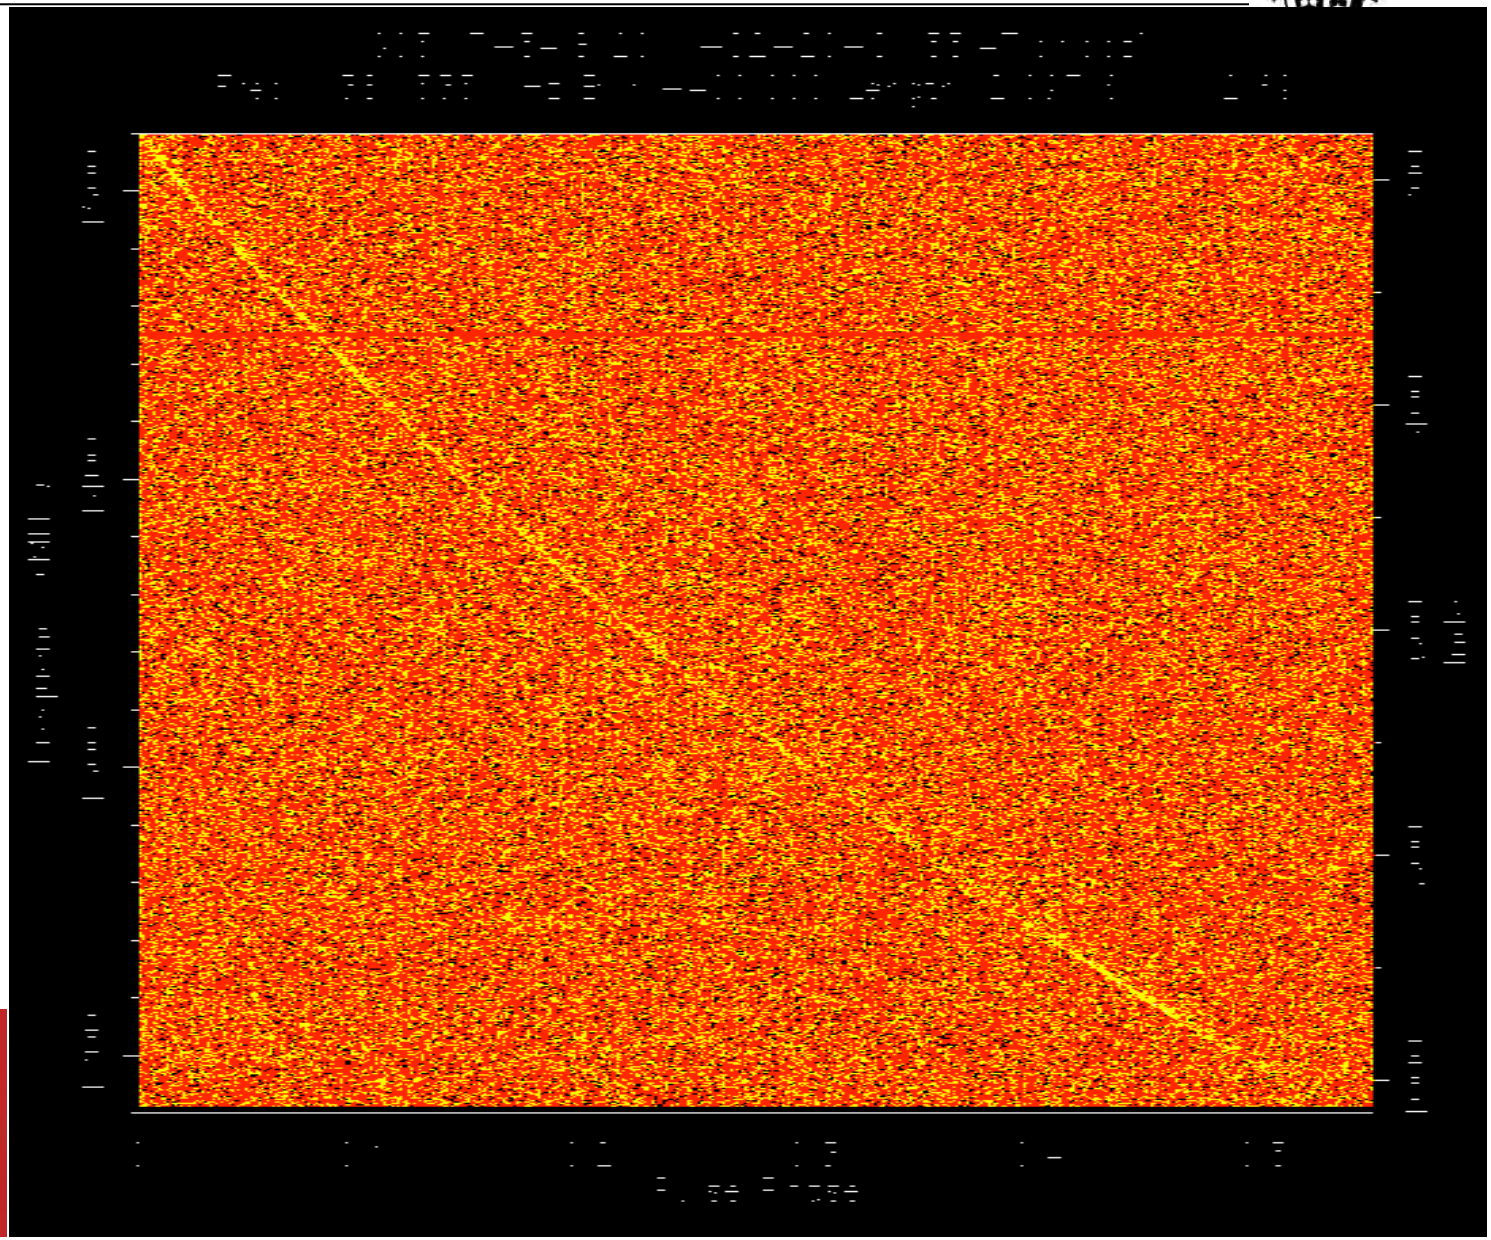

SWINBURNE UNIVERSITY OF TECHNOLOGY

More bursts

- DMs
  - 995 pc/cc
  - 723 pc/cc
  - 1103 pc/cc
  - 553 pc/cc
- + lots more!

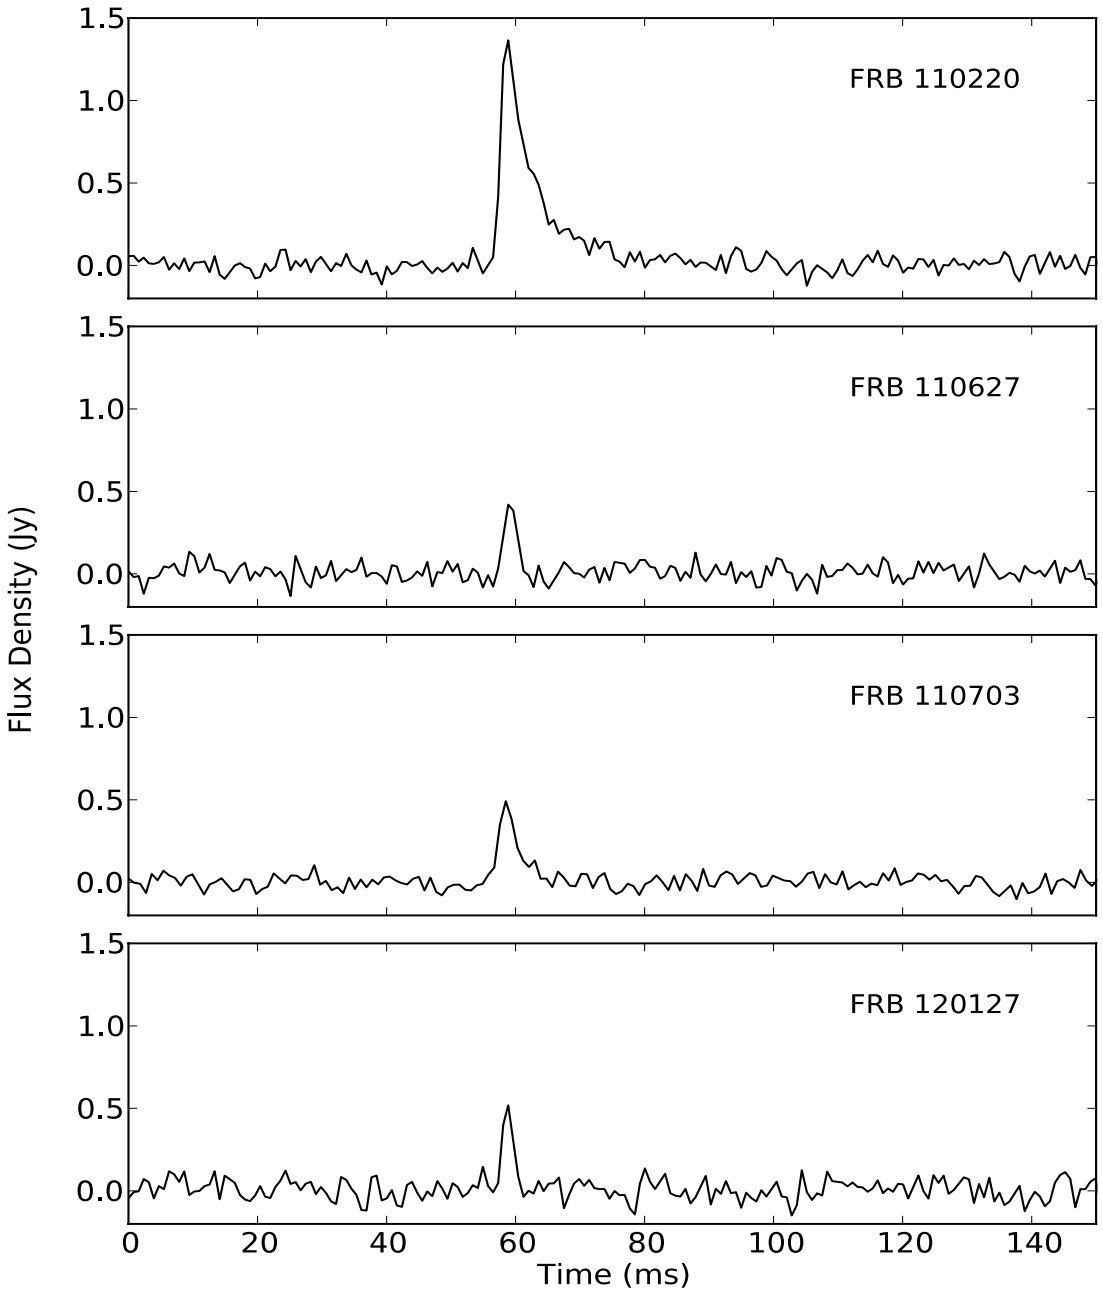

SWINBURNE UNIVERSITY OF TECHNOLOGY

Matthew Bailes + HITRUN team.

A Population of Fast Radio Bursts at Cosmological Distances

D. Thornton,[1,2*] B. Stappers,[1] M. Bailes,[3,4] B. Barsdell,[3,4] S. Bates,[5] N. D. R. Bhat,[3,4,6] M. Burgay,[7] S. Burke-Spolaor,[8] D. Champion,[9] P. Coster,[2,3] N. D'Amico,[7,10] A. Jameson,[3,4] S. Johnston,[2] M. Keith,[2] M. Kramer,[9,1] L. Levin,[5] S. Milia,[7] C. Ng,[9] A. Possenti,[7] W. van Straten[3,4]

Searches for transient astrophysical sources often reveal unexpected classes of objects that are useful physical laboratories. In a recent survey for pulsars and fast transients, we have uncovered four millisecond-duration radio transients all more than 40° from the Galactic plane. The bursts' properties indicate that they are of celestial rather than terrestrial origin. Host galaxy and intergalactic medium models suggest that they have cosmological redshifts of 0.5 to 1 and distances of up to 3 gigaparsecs. No temporally coincident e- or gamma-ray signature was identified in association with the bursts. Characterization of the source population and identification of host galaxies offers an opportunity to determine the baryonic content of the universe.

SWINBURNE UNIVERSITY OF TECHNOLOGY

CAASTRO
ARC CENTRE OF EXCELLENCE FOR ALL-SKY ASTROPHYSICS

Where we've looked, where we've found

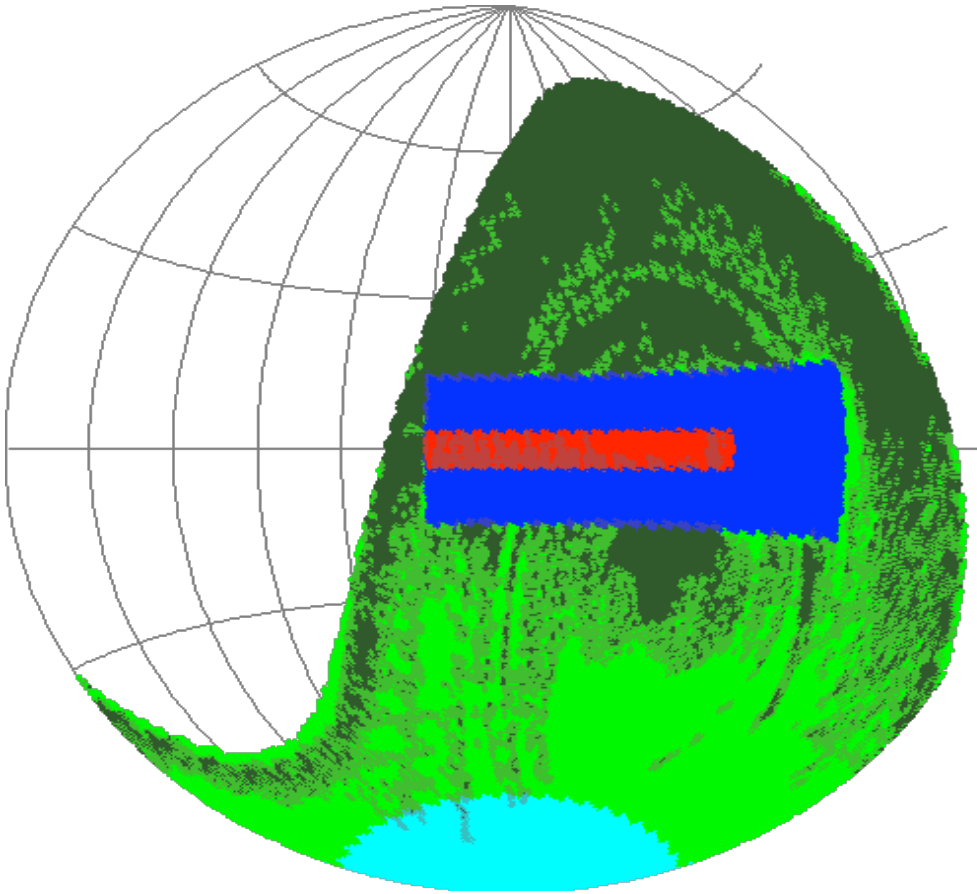

(bit out of date)

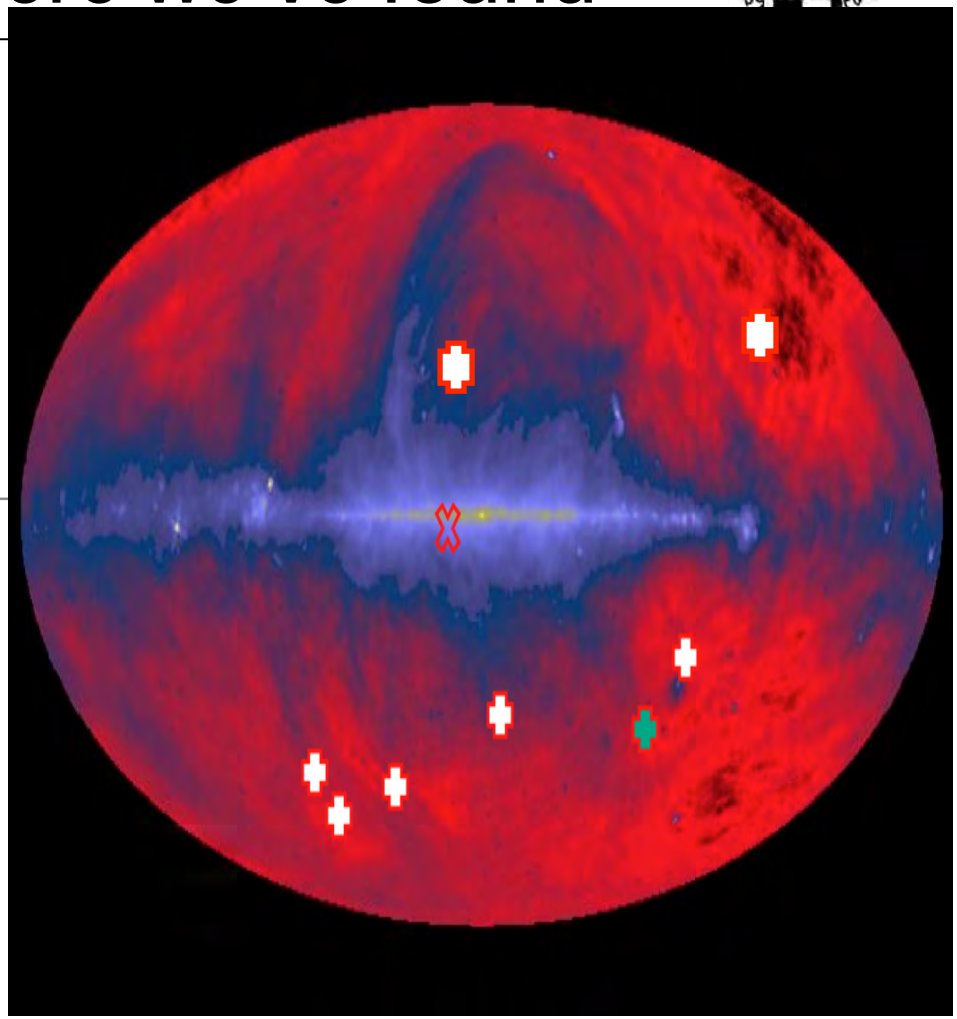

If RFI, should see just as many on the plane?

HITRUN Lorimer Burst Rates

- Implied rates: many thousand/day/sky
  - Too high for ns+ns (SNe?)
- Ours have Z ~0.5-1
- Free electron count of the Universe?
- DM delay to MWA is tens of seconds!
  - Catch one at another telescope?

SWINBURNE UNIVERSITY OF TECHNOLOGY

Upgrade #1 real-time detector!

- Uses GPUs to dedisperse and search in real time.
- Ben Barsdell PhD thesis

SWINBURNE UNIVERSITY OF TECHNOLOGY

Barsdell's Heimdall Real-time RRAT Detector

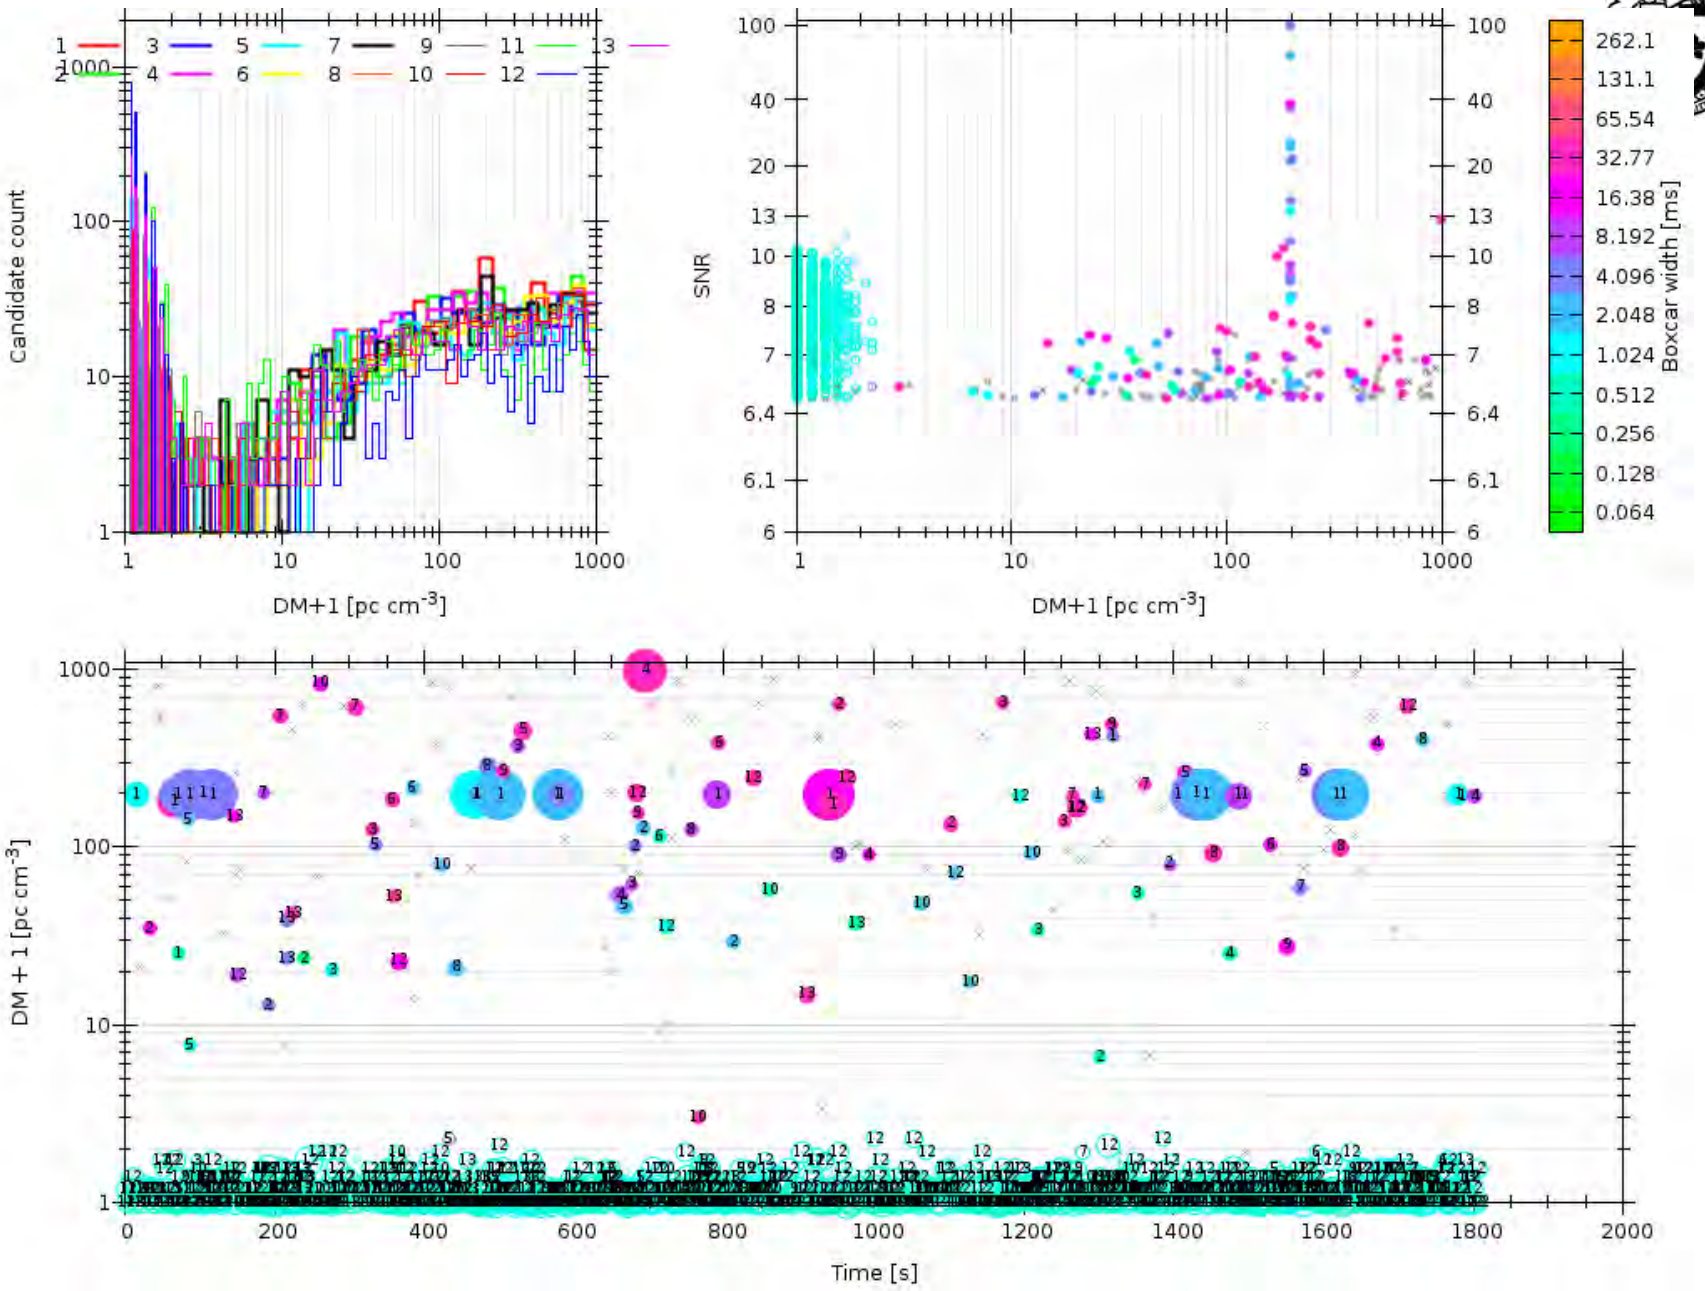

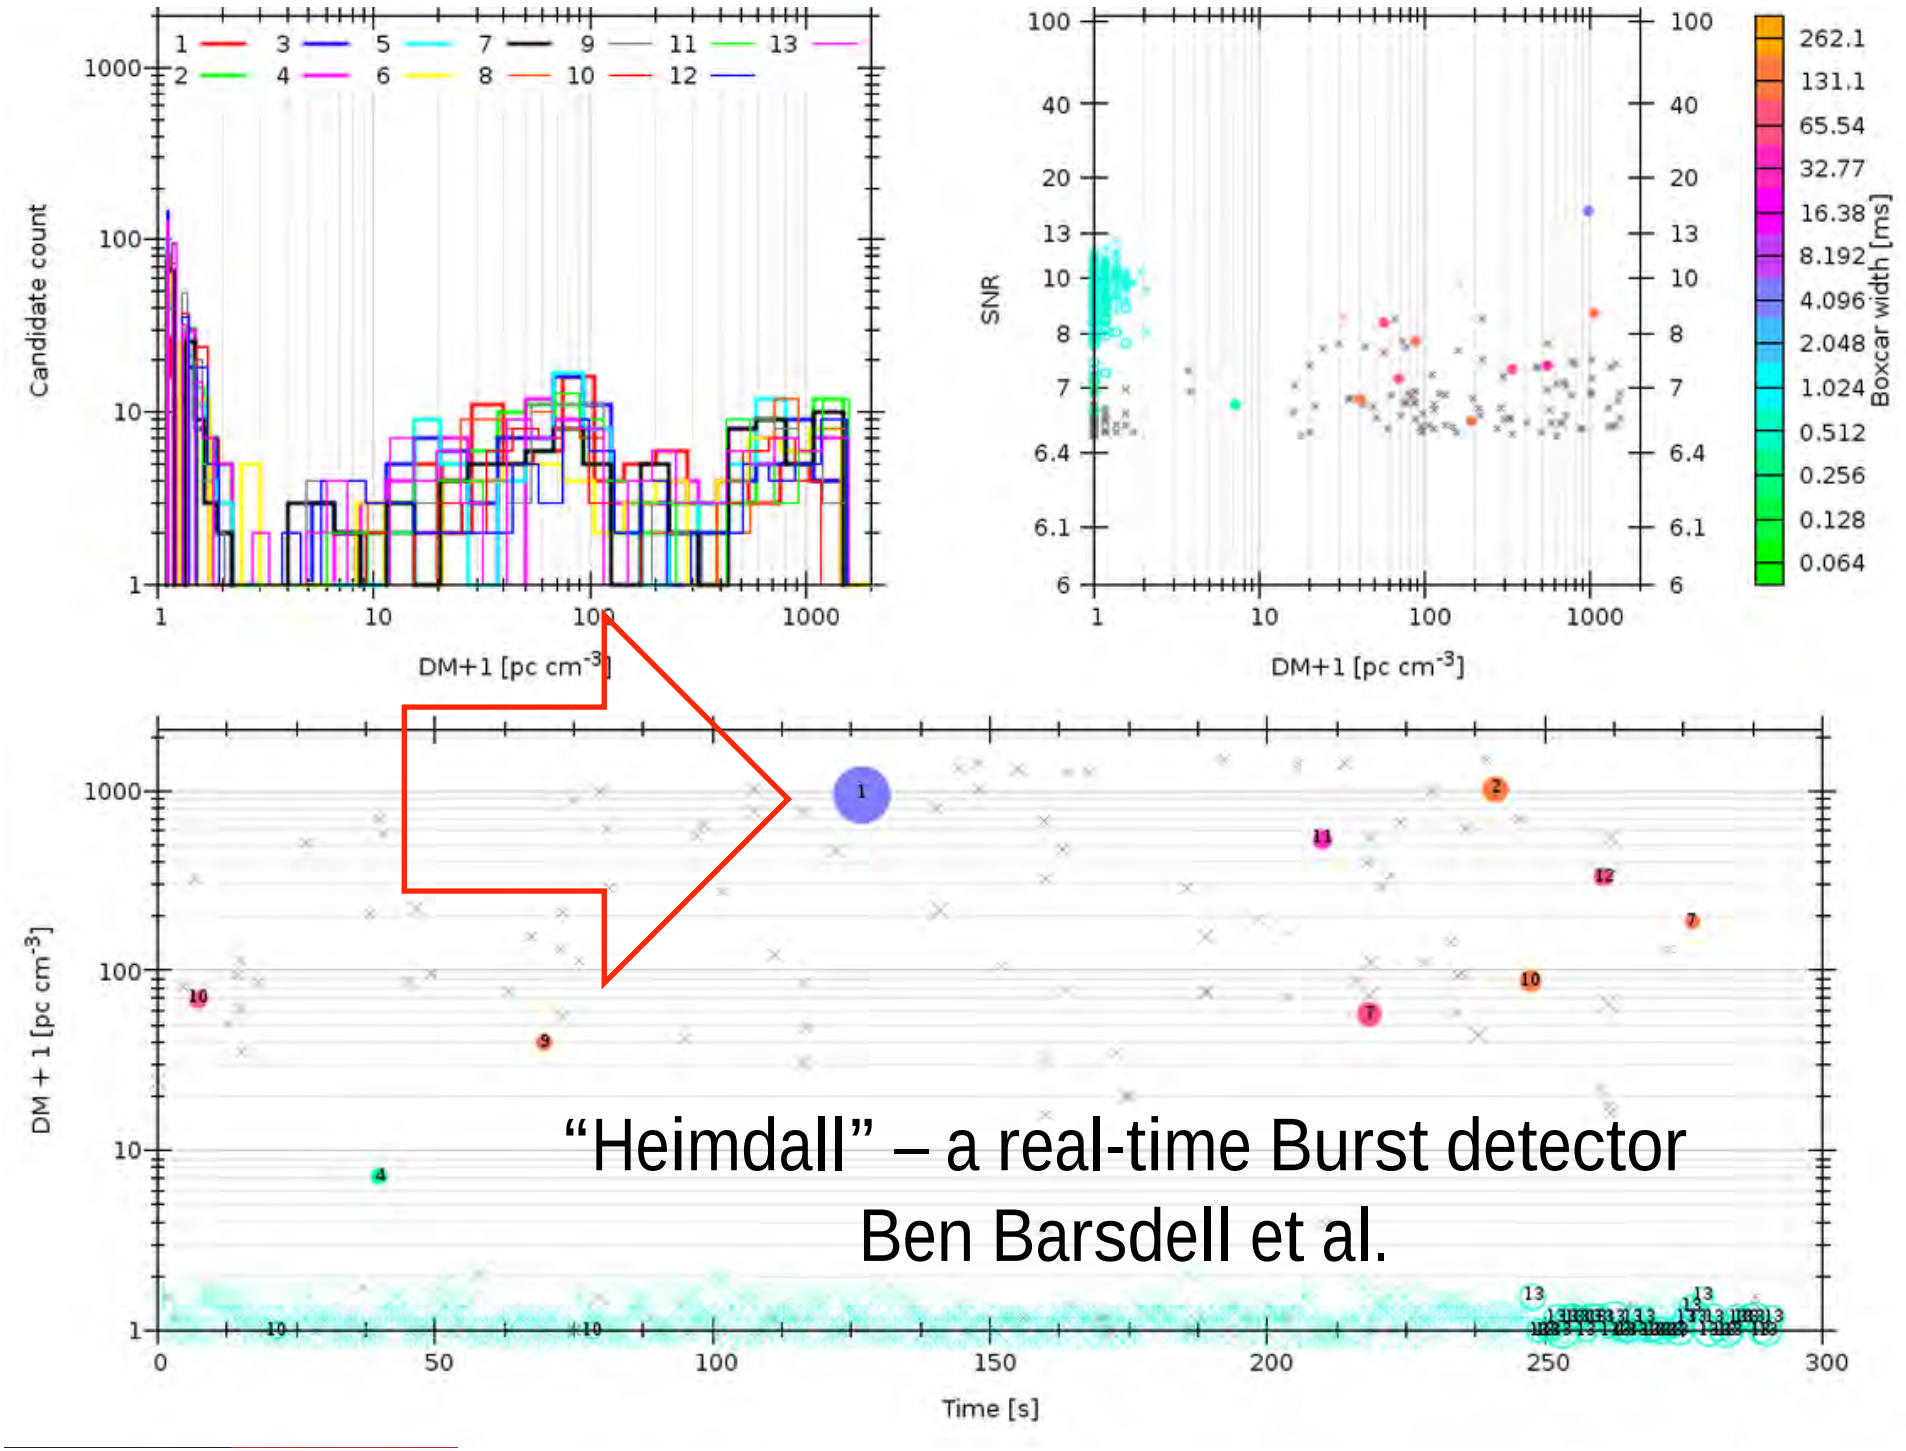
Candidate count
DM+1 [pc cm^{-3}]
SNR
Boxcar width [ms]
DM + 1 [pc cm^{-3}]
Time [s]
"Heimdall" – a real-time Burst detector
Ben Barsdell et al.

SkyMapper Trigger?

- “know” about an FRB within a few seconds.

SWINBURNE UNIVERSITY OF TECHNOLOGY

Molonglo Radio Telescope

Anne Green: Telescope Director

Duncan Campbell-Wilson →

Dick Hunstead (Sydney)

Tim Bateman (CASS)

Russ McWhirter (Haystack)

Andrew Jameson (Swinburne)

Matthew Bailes, Ewan Barr, Chris Flynn, Evan Keane, Fabian Jankowski (Swinburne)

Manisha Caleb (ANU), Nie Jun (Urumqi)

Molonglo

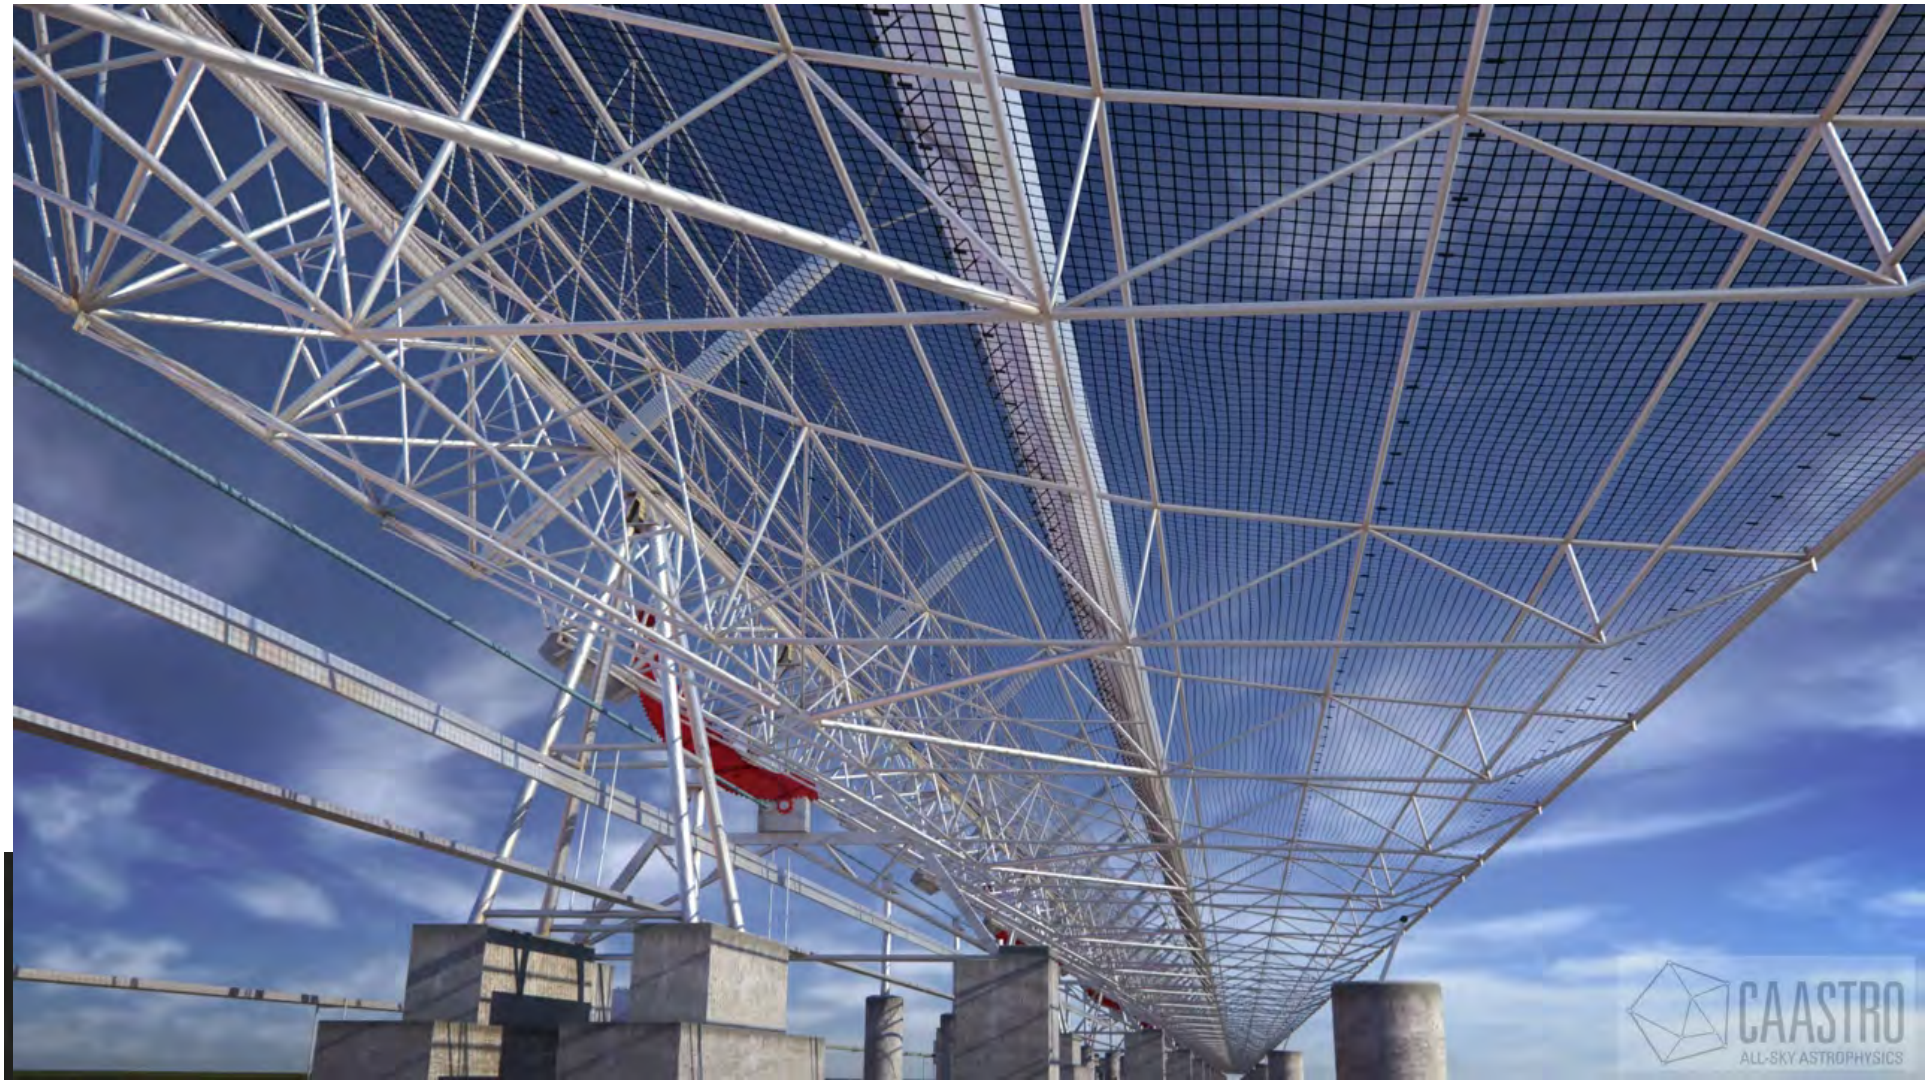

Background

- MOST is the largest radio telescope in Australia: 5xPKS, 4xASKAP
- Development work on digital back-end began around 2005: PFB + correlator, 100 MHz BW, 700-1100 MHz
  - Project backend stalled.
- August 2012 Swinburne floated alternative correlator solution: GPU cluster, with fast sampling (10 Gb/s, 10 GbE)

SWINBURNE UNIVERSITY OF TECHNOLOGY

“Add by 22 = Phased array Beam”

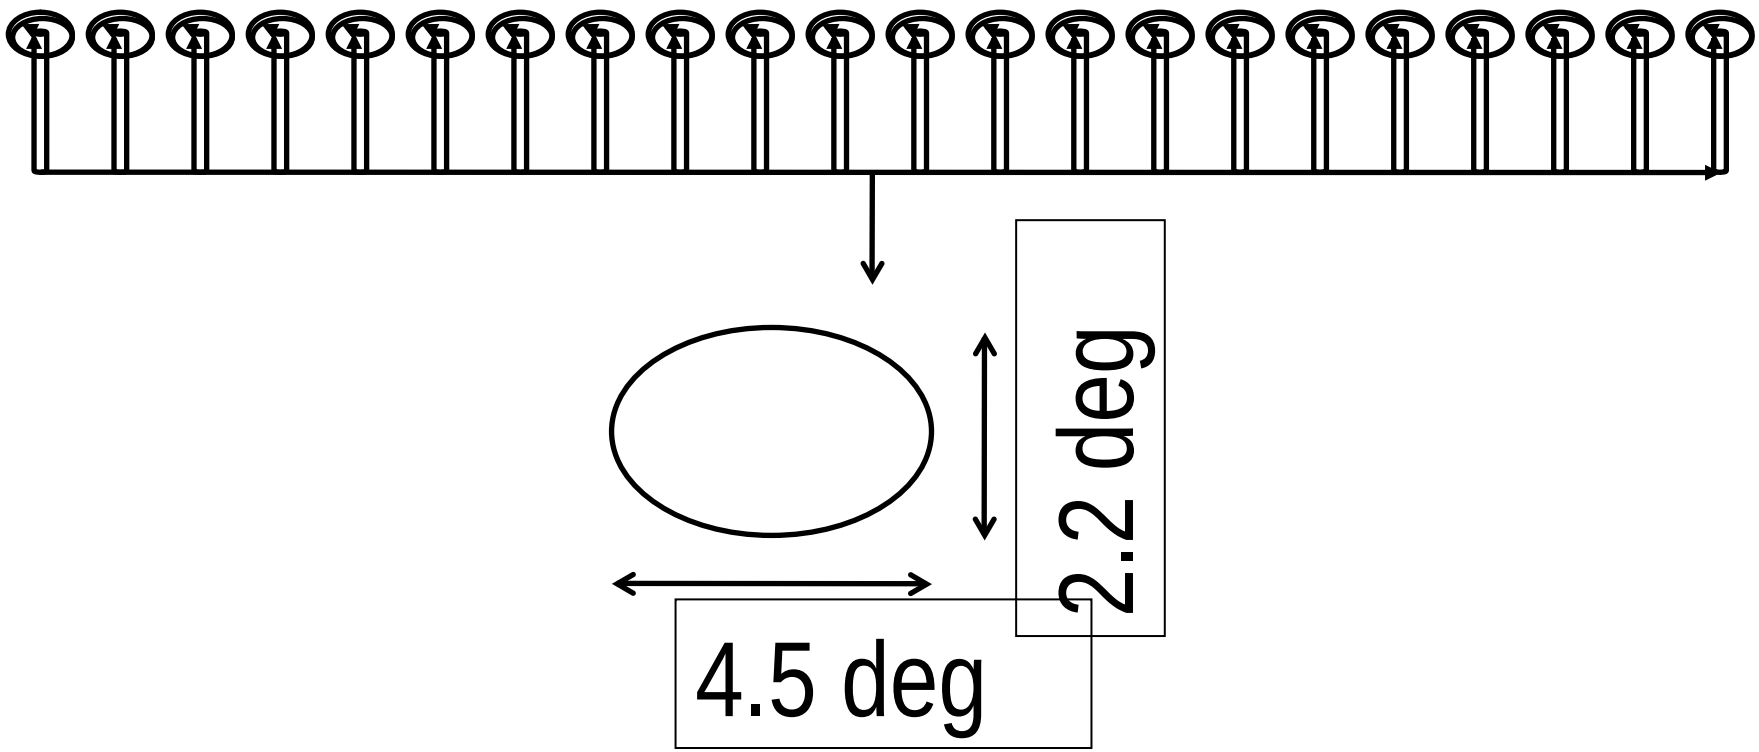

100 MB/s * 352

SWINBURNE UNIVERSITY OF TECHNOLOGY

New CPU/GPU Design (dual 10Gb NIC)

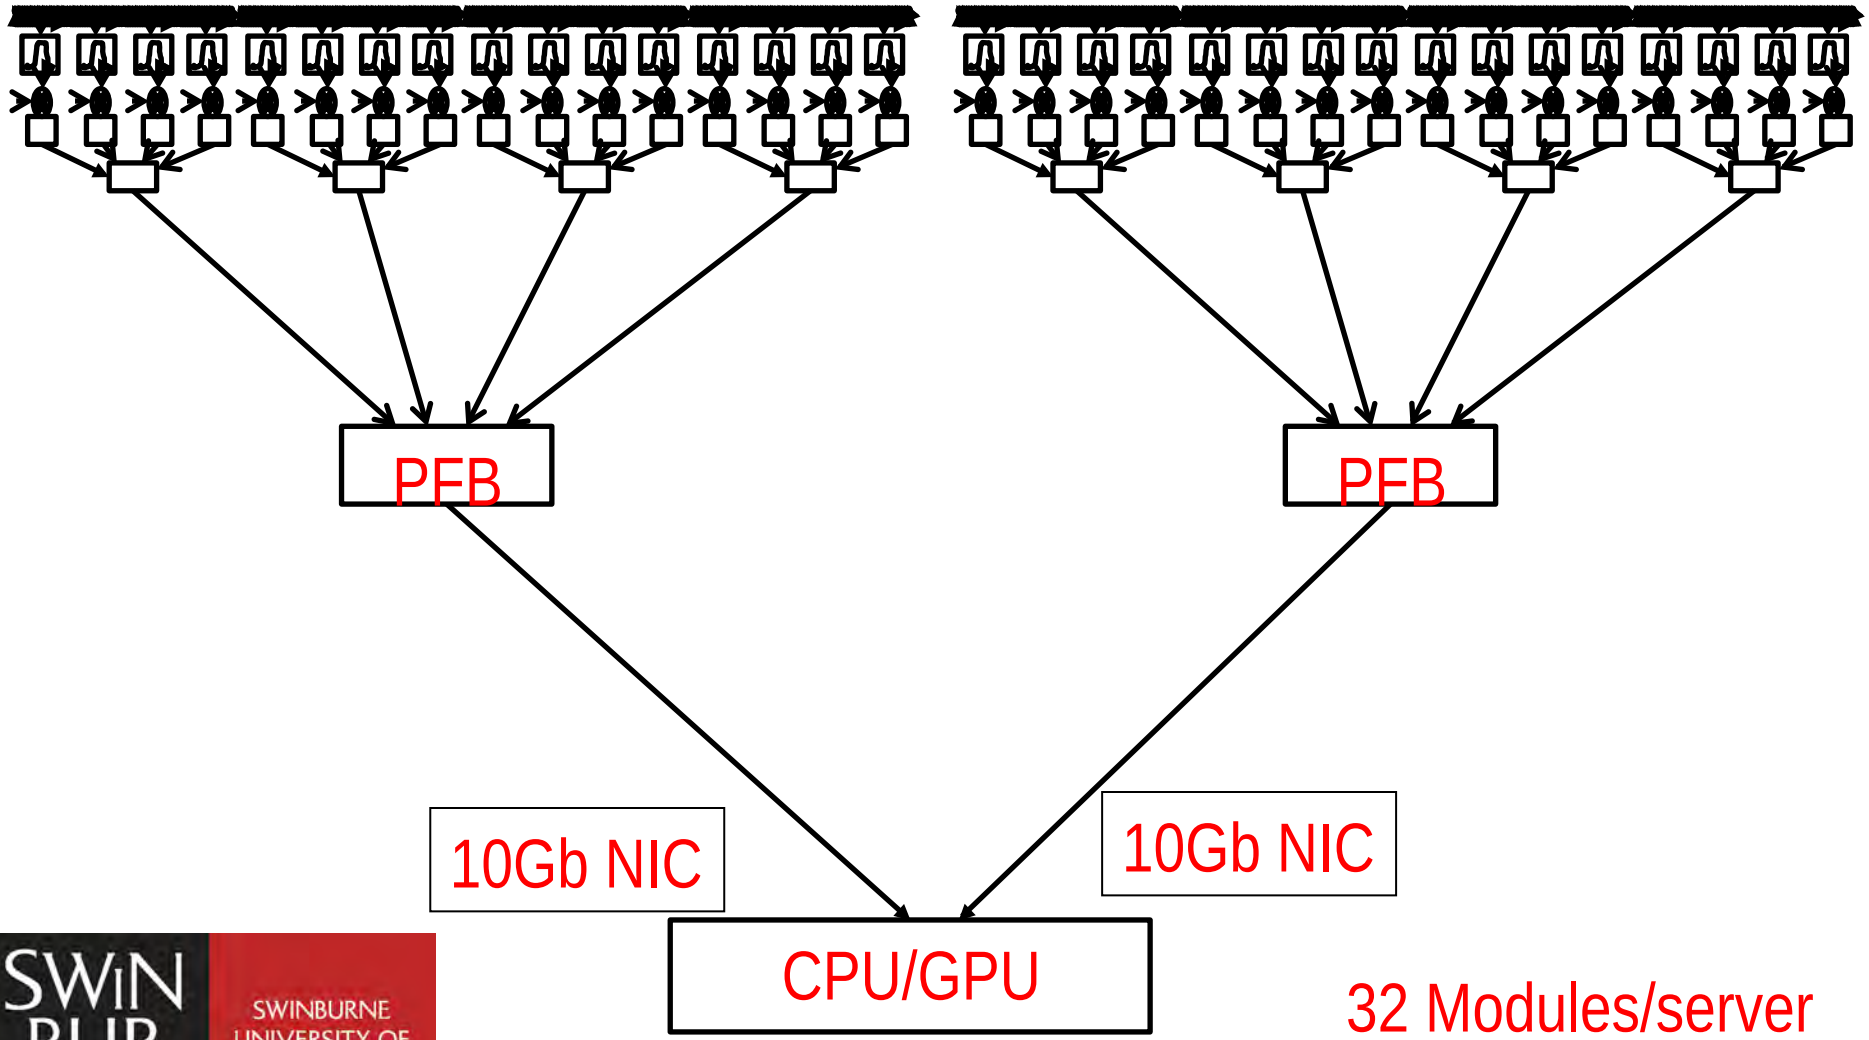

32 Modules/server

SWINBURNE UNIVERSITY OF TECHNOLOGY

Data Rates

- 30 MHz * 1 poln * 8 bits * 2 (Nyquist) * 352 * 86400 * 365 = 666 Petabytes/year!

SWINBURNE UNIVERSITY OF TECHNOLOGY

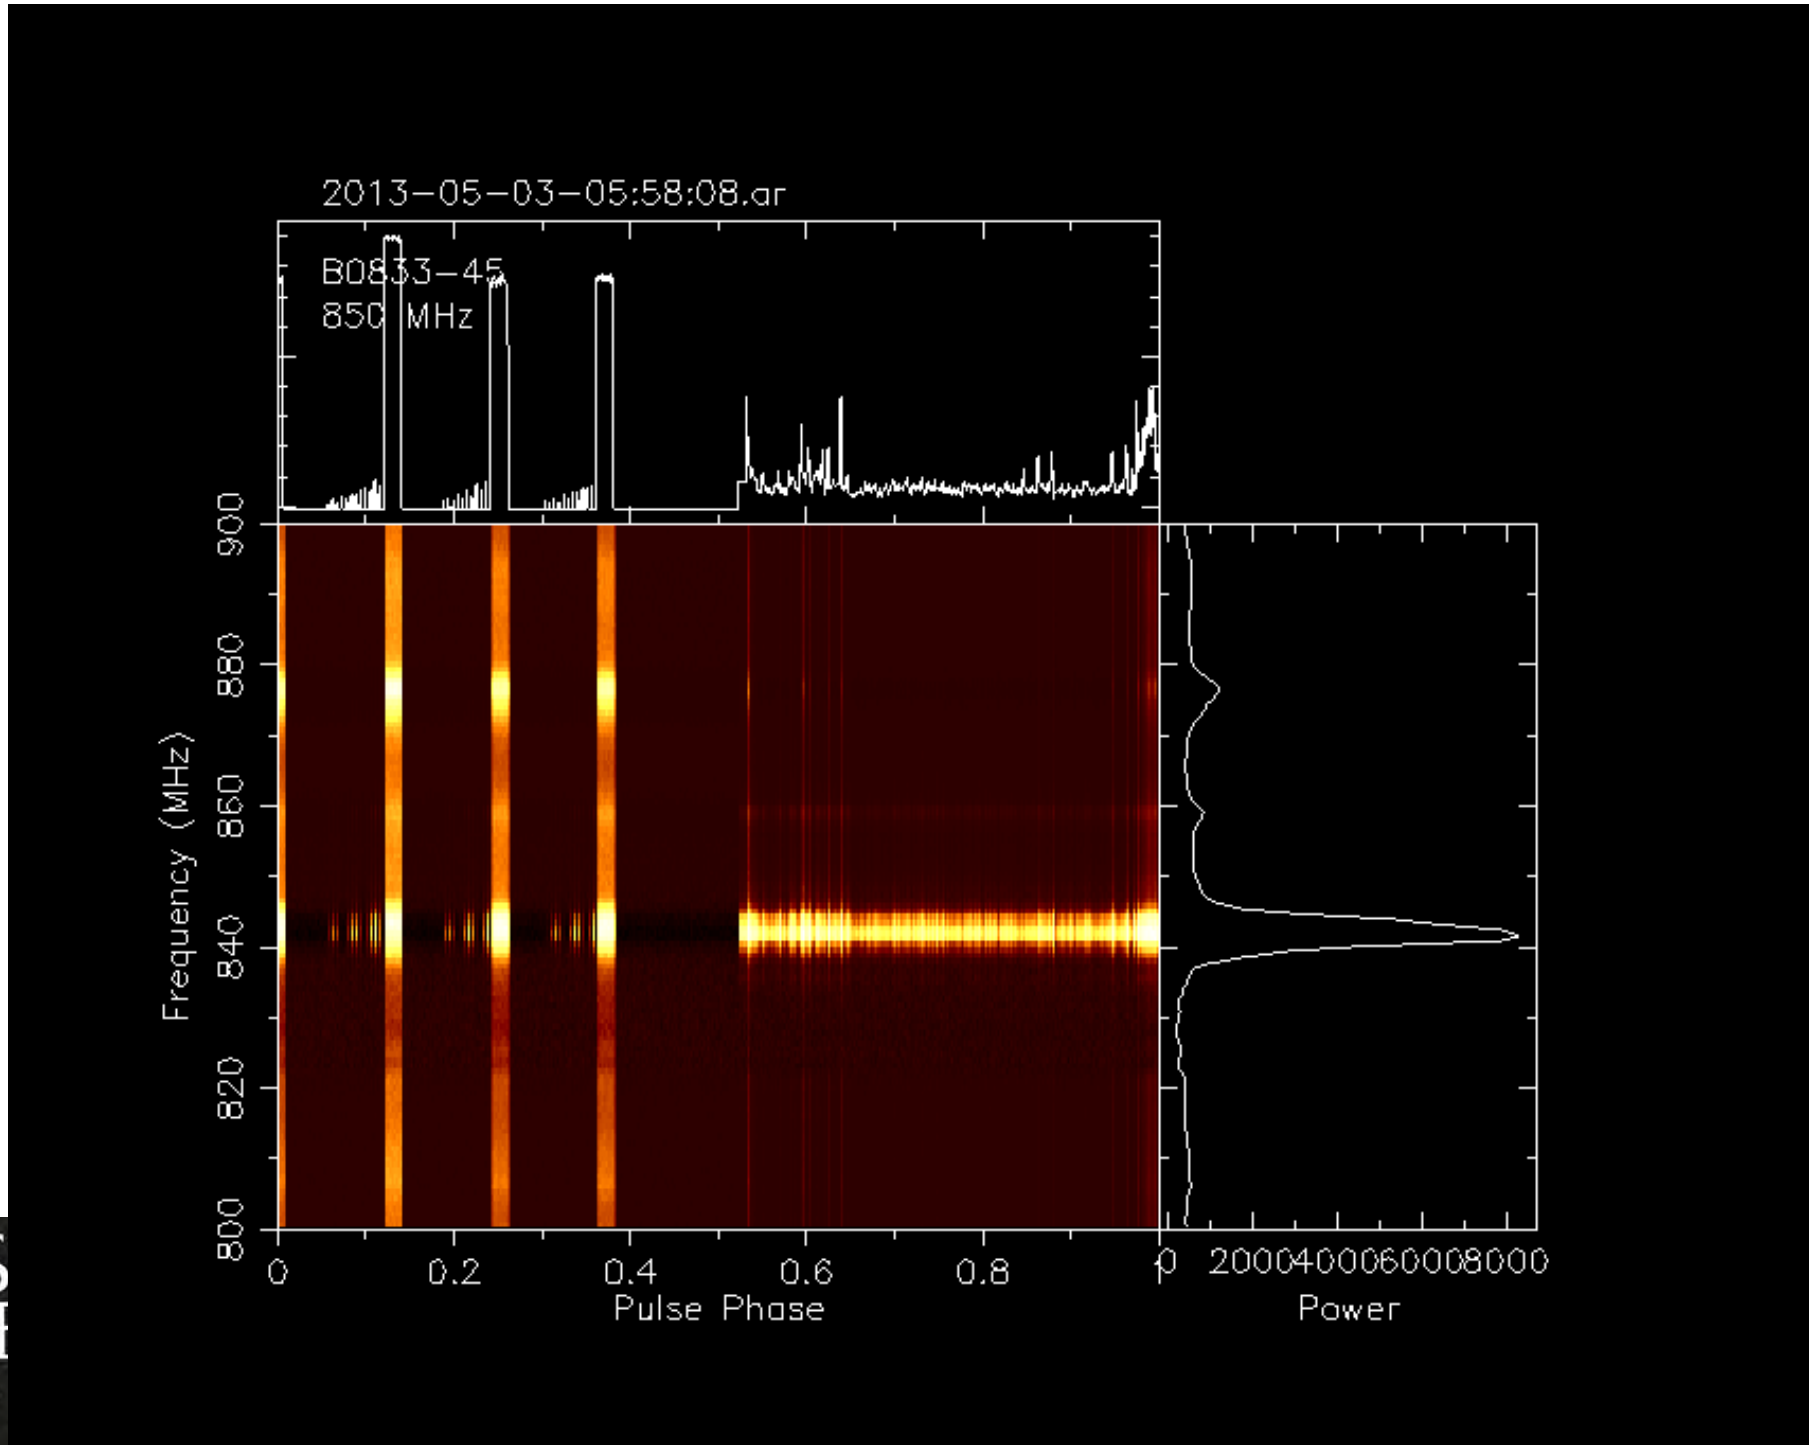
2013-05-03-05:58:08.ar
B0833-45
850 MHz
Frequency (MHz)
800
820
840
860
880
900
0
0.2
0.4
0.6
0.8
Pulse Phase
0
2000
4000
6000
8000
Power

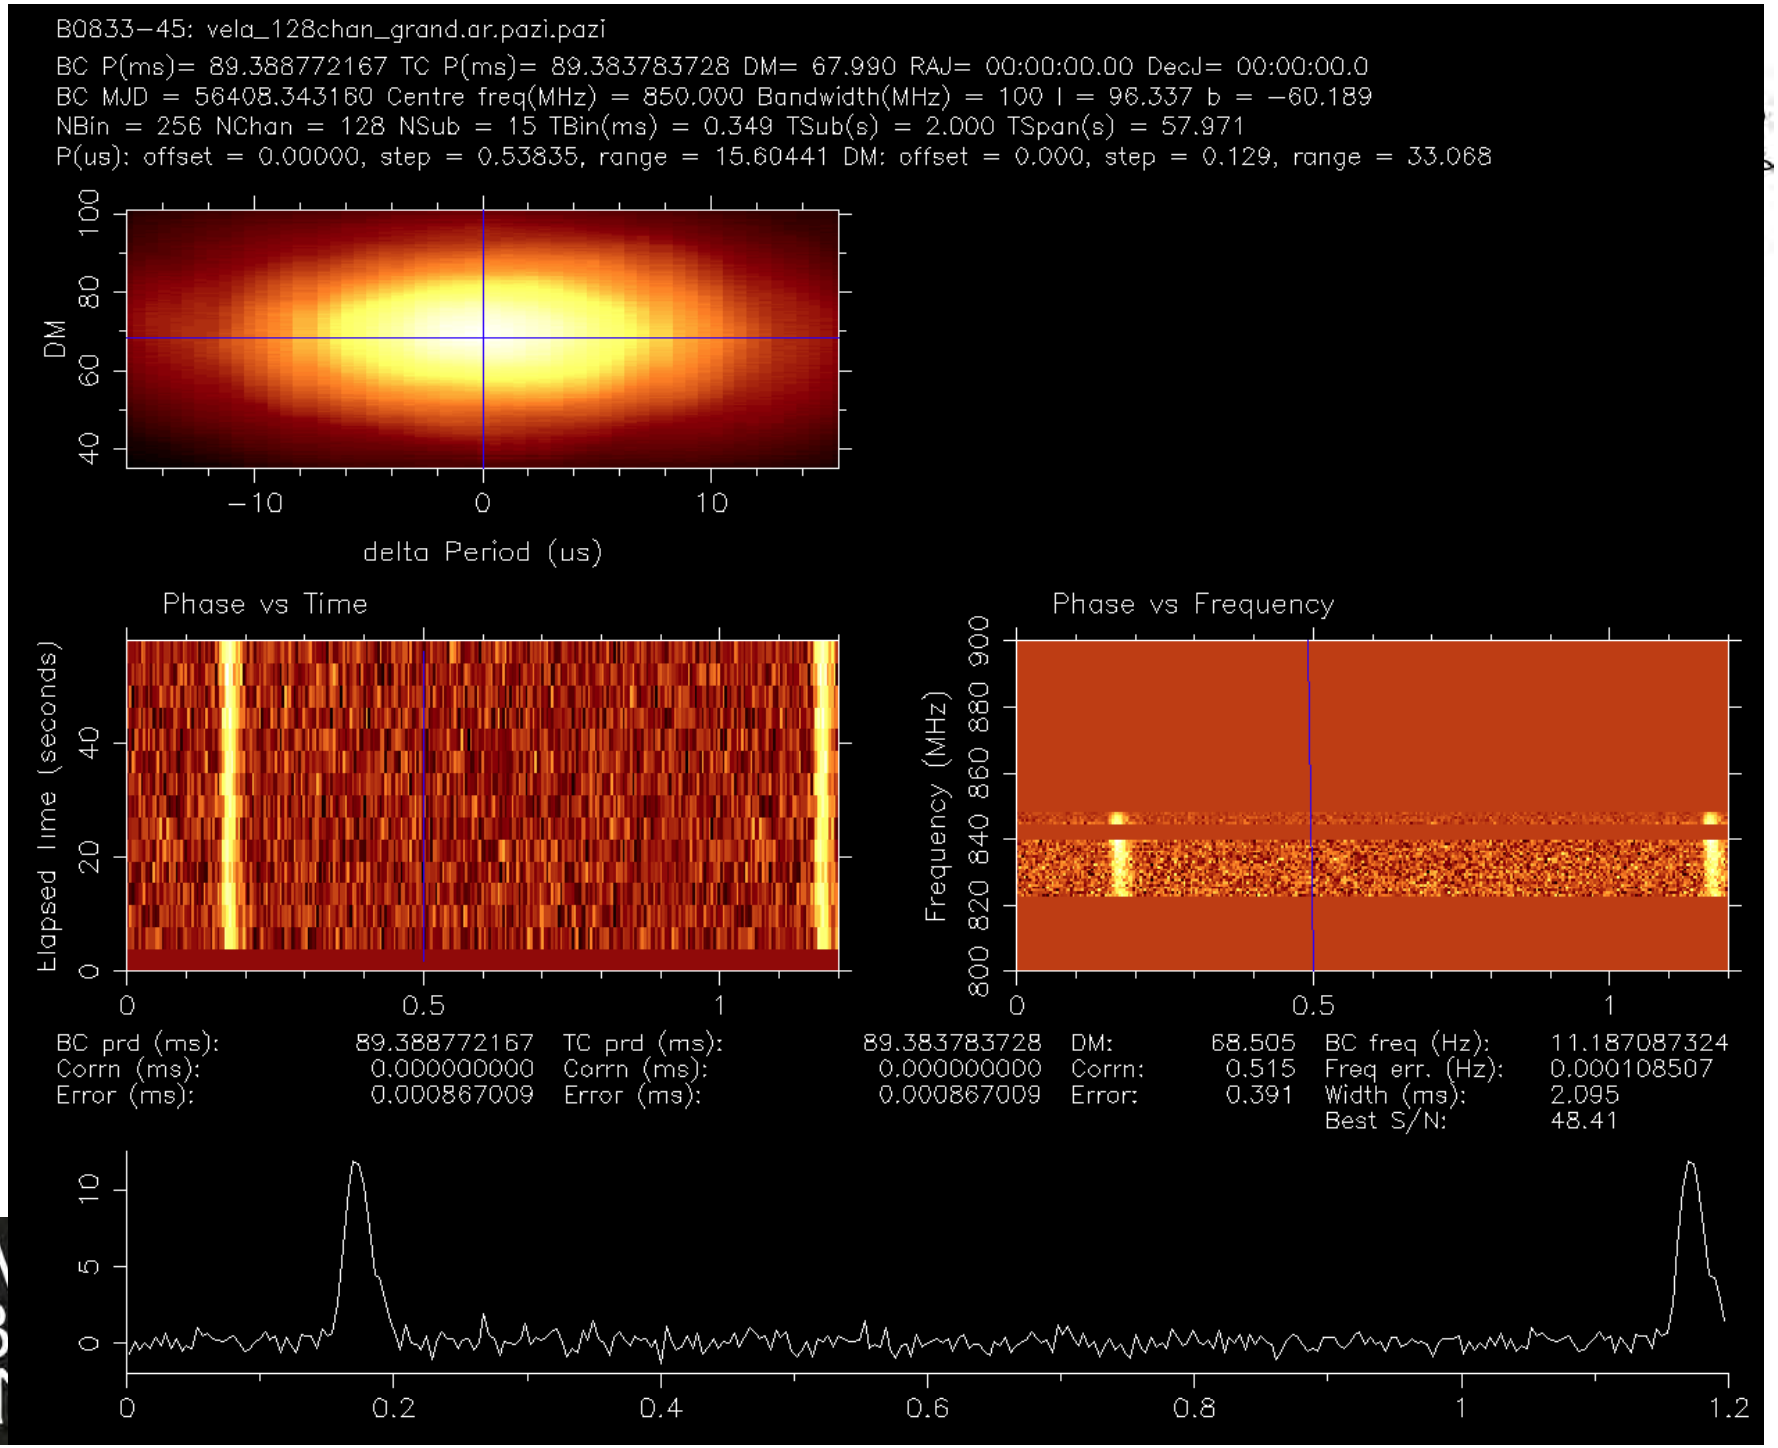
B0833−45: vela_128chan_grand.ar.pazi.pazi
BC P(ms)= 89.388772167 TC P(ms)= 89.383783728 DM= 67.990 RAJ= 00:00:00.00 DecJ= 00:00:00.0
BC MJD = 56408.343160 Centre freq(MHz) = 850.000 Bandwidth(MHz) = 100 l = 96.337 b = −60.189
NBin = 256 NChan = 128 NSub = 15 TBin(ms) = 0.349 TSub(s) = 2.000 TSpan(s) = 57.971
P(us): offset = 0.00000, step = 0.53835, range = 15.60441 DM: offset = 0.000, step = 0.129, range = 33.068
DM
delta Period (us)
Phase vs Time
Elapsed Time (seconds)
Phase vs Frequency
Frequency (MHz)
BC prd (ms): 89.388772167
Corrn (ms): 0.000000000
Error (ms): 0.000867009
TC prd (ms): 89.383783728
Corrn (ms): 0.000000000
Error (ms): 0.000867009
DM: 68.505
Corrn: 0.515
Error: 0.391
BC freq (Hz): 11.187087324
Freq err. (Hz): 0.000108507
Width (ms): 2.095
Best S/N: 48.41

Millisecond pulsar (Johnston et al 1993)!

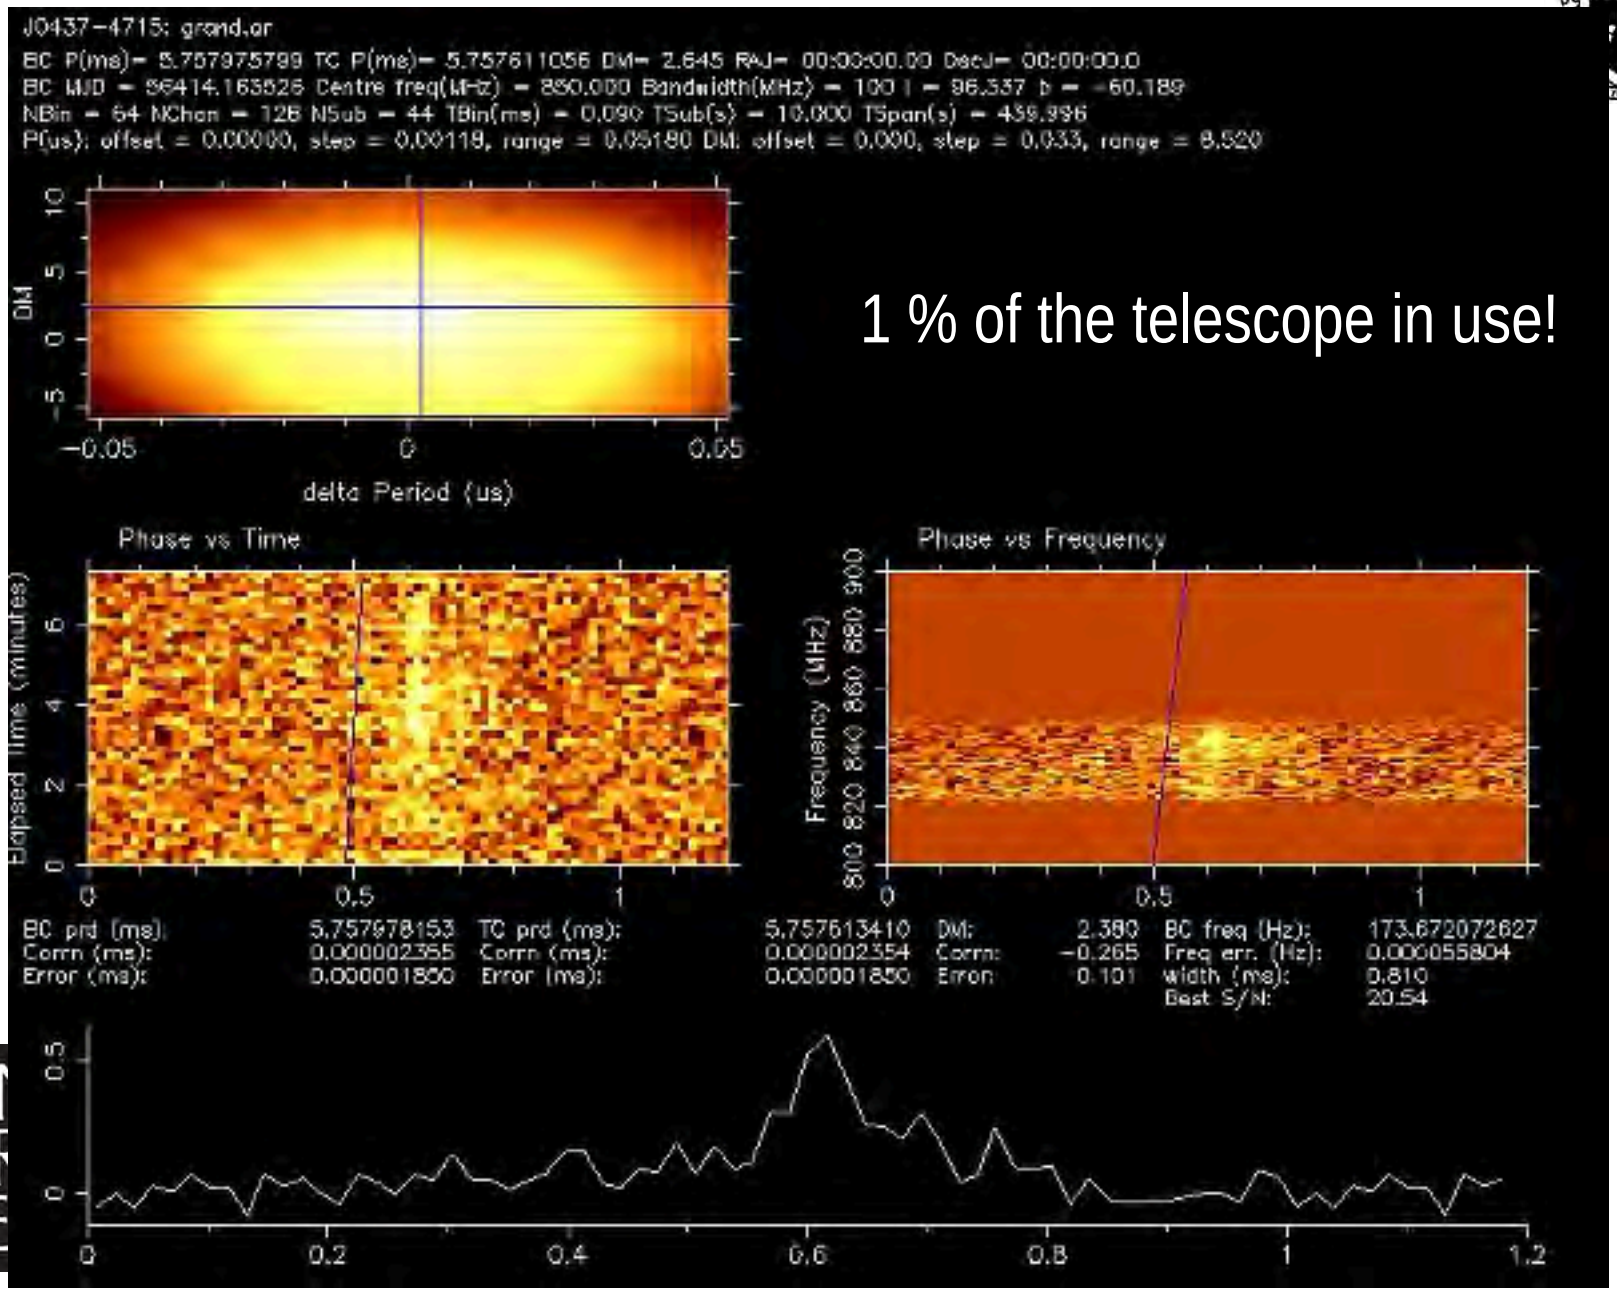

1 % of the telescope in use!

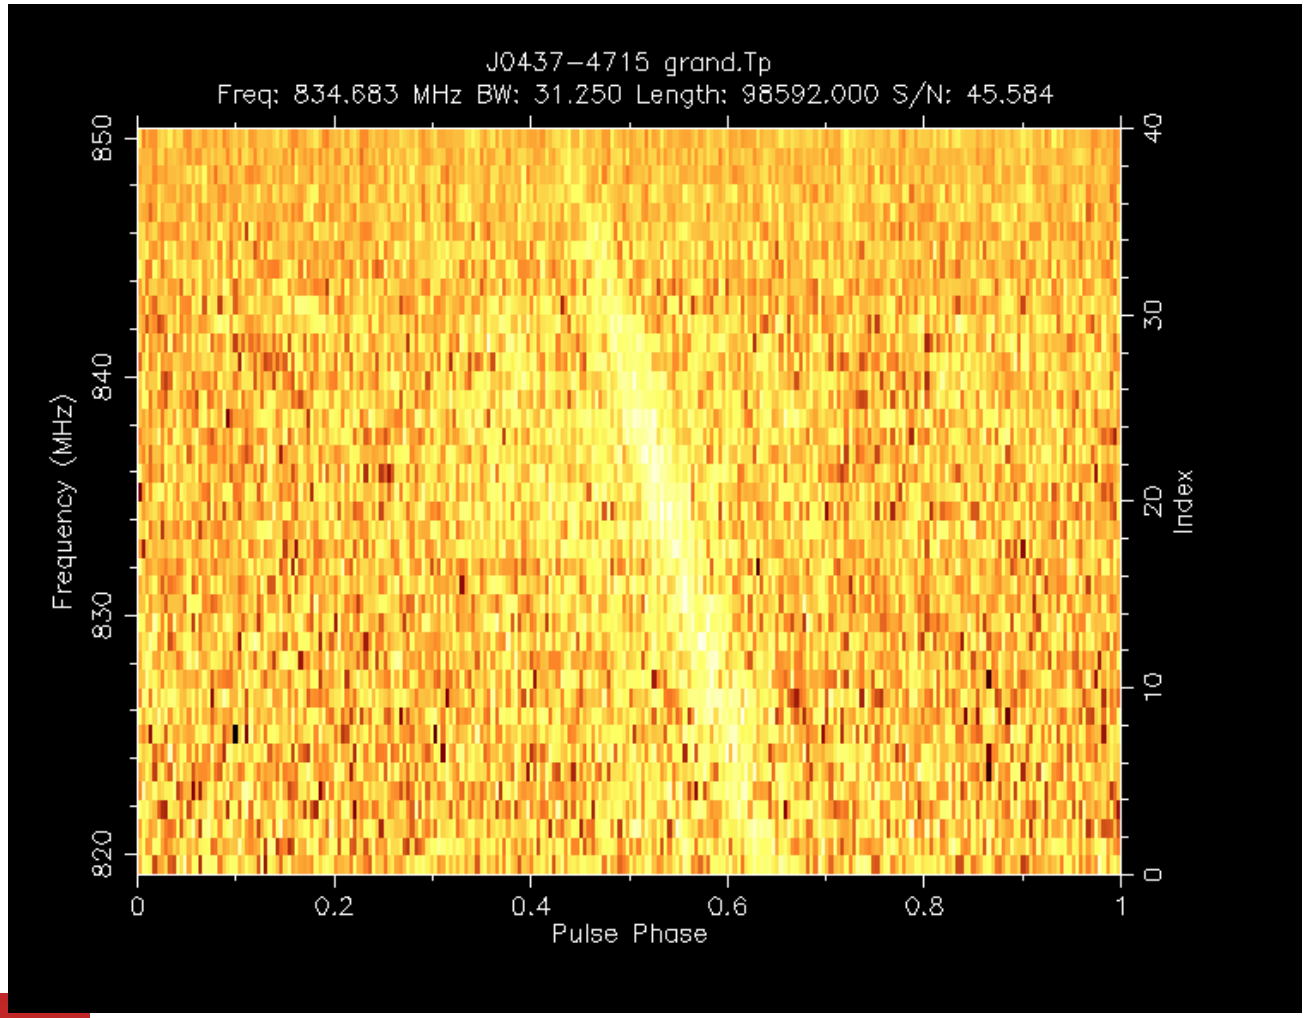

SWINBURNE UNIVERSITY OF TECHNOLOGY

New Design (dual 10Gb NIC)

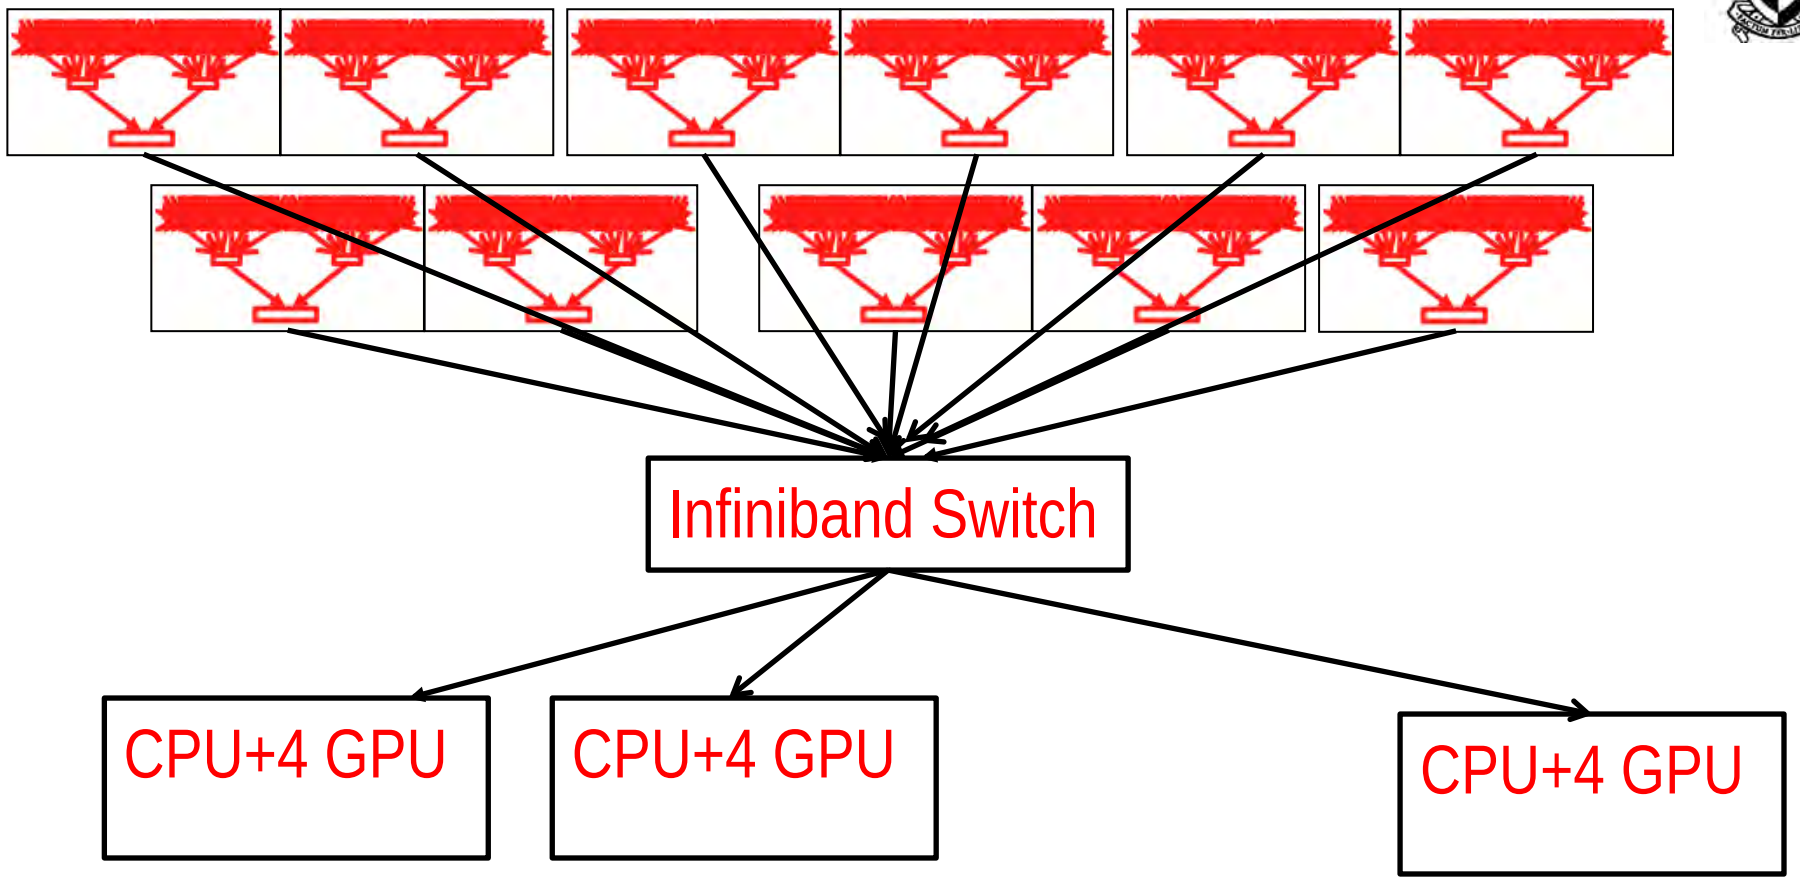

6 quad-GPU servers

SWINBURNE UNIVERSITY OF TECHNOLOGY

Interesting bits

- Each server catches from its Network card:
  - 30 x 2 * 22 = 1.92 GB/s! (x11)
- Sent to GPU
  - 1.92 GB/s (x11)
- Expanded to 32 bits
  - 7.68 GB/s (x11)
  - Gaussian? (Y = OK, N = mask)
- Returned to RAM after 8-bitting
  - 1.92 GB/s (x11)

SWINBURNE UNIVERSITY OF TECHNOLOGY

Corner Turn

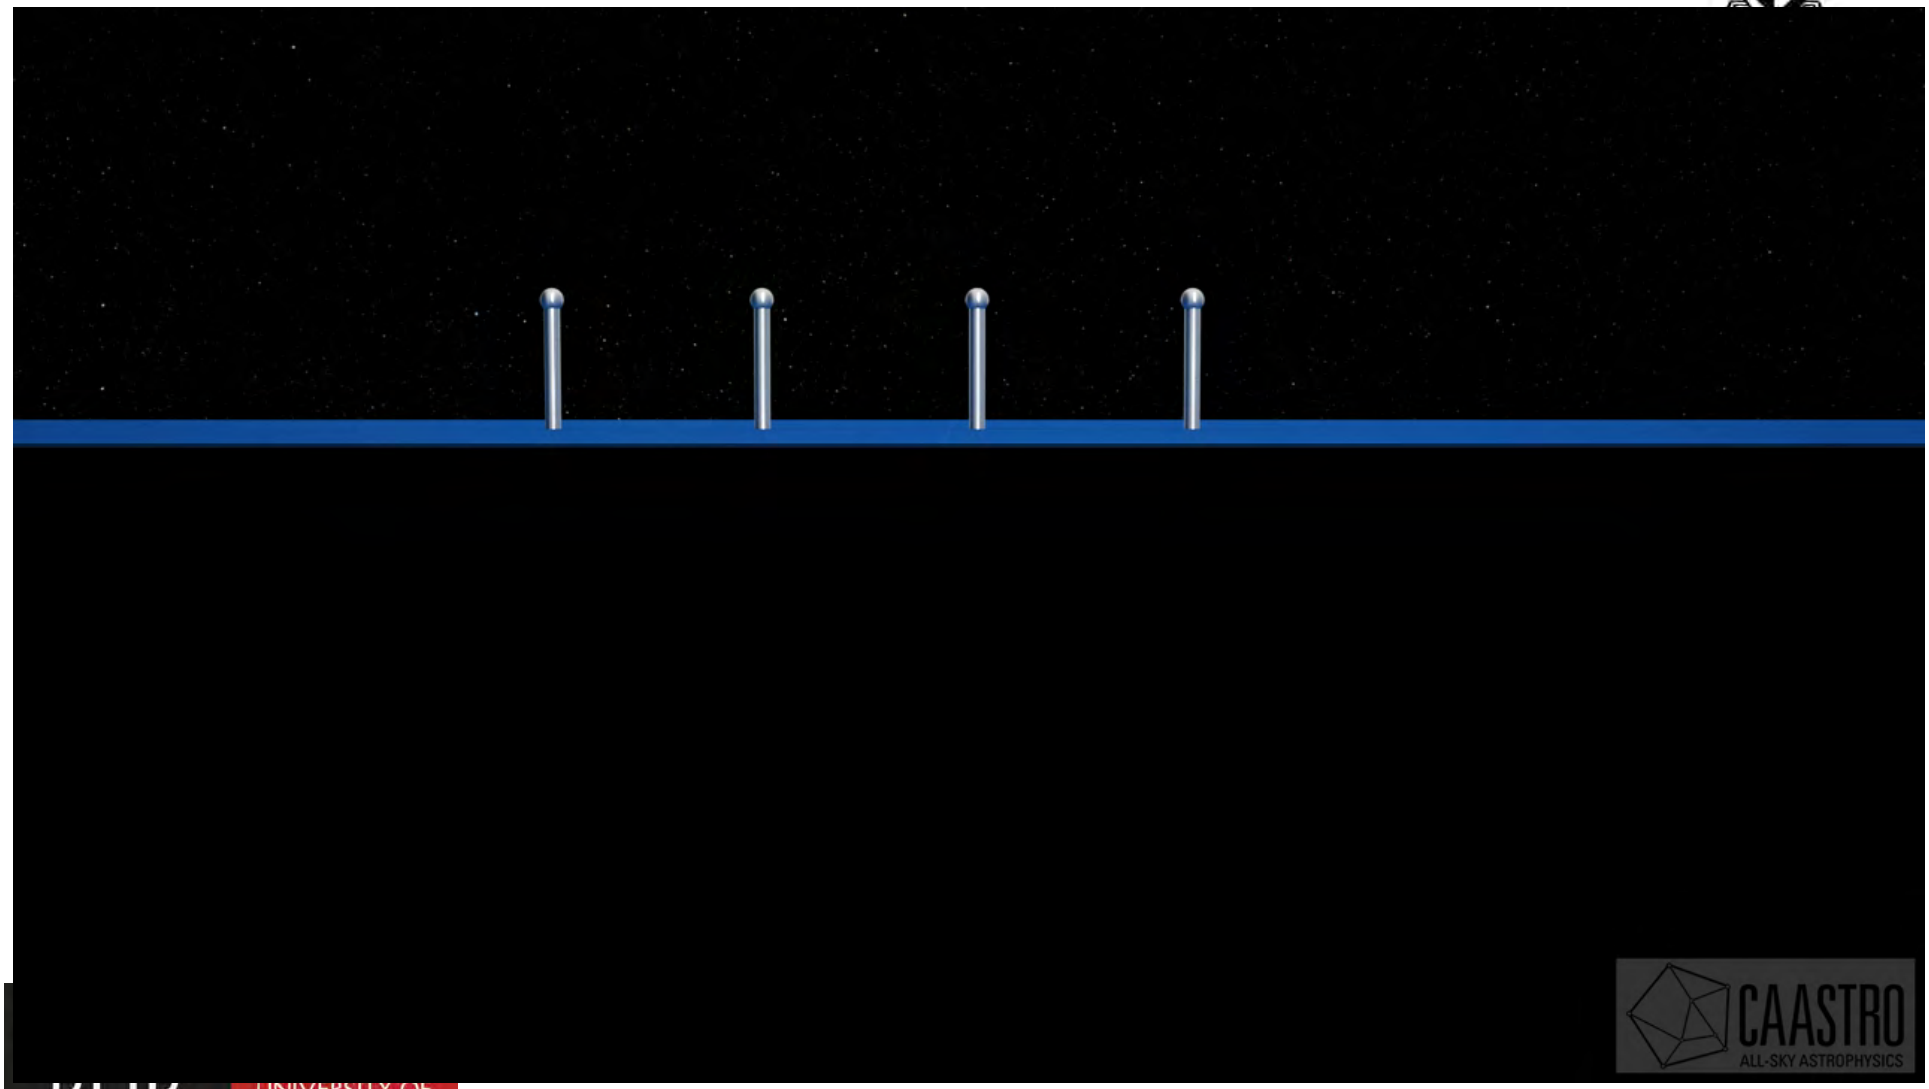

RDMA over IB

- IB is “infiniband” high speed low latency interconnect
- 10K 56 Gb/s switch
  - $300/NIC
  - $200/cable
- Can do ~4.4 GB/s without using the CPU using RDMA.
- Back to GPUs
  - Spatial FFT (fan beams)

SWINBURNE UNIVERSITY OF TECHNOLOGY

Interesting bits

- GPUs for RFI excision & incoherent folding
- RDMA over IB to do the corner turn (no CPU)
  - 4445 MB/s = 35.6 Gb/s
- 2D FFT to create 512 fan beams + finer channels
  - Whole instrument on 7 x 690s
- Raw dump mode (30 s)
  - RFI “parallax”

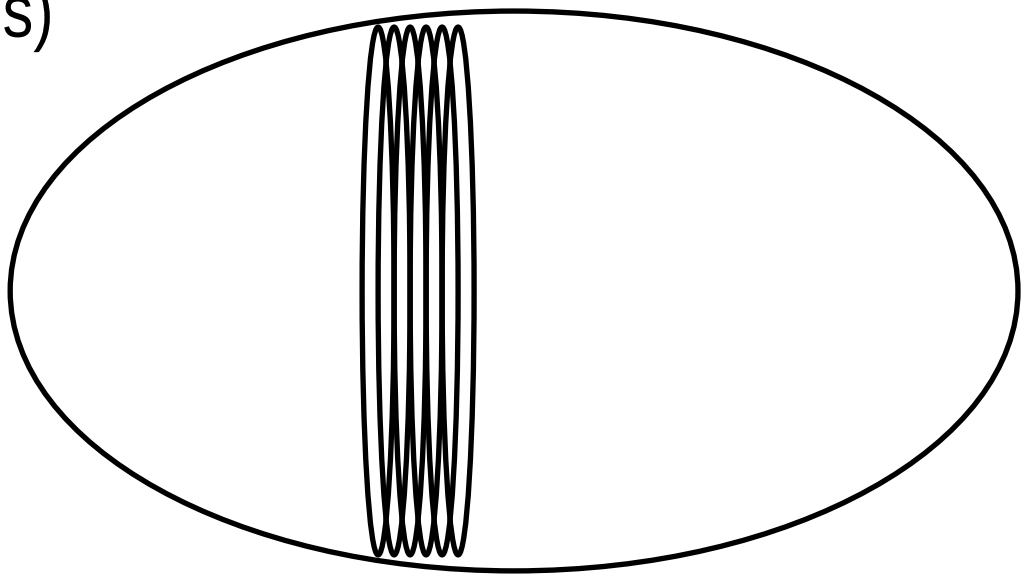

SWINBURNE UNIVERSITY OF TECHNOLOGY

“Burst” mode

- Search fan beams for Lorimer bursts (FRBs)
- GPU dedisperser + “Heimdall”
  - Barsdell’s PhD thesis.
- Also find RRATs, single pulses from pulsars
- (Should) find 1 burst per few days > 10 sigma
- Real-time detector
- Raw dump from ring buffer
- Position 43”/SNR & 2 deg

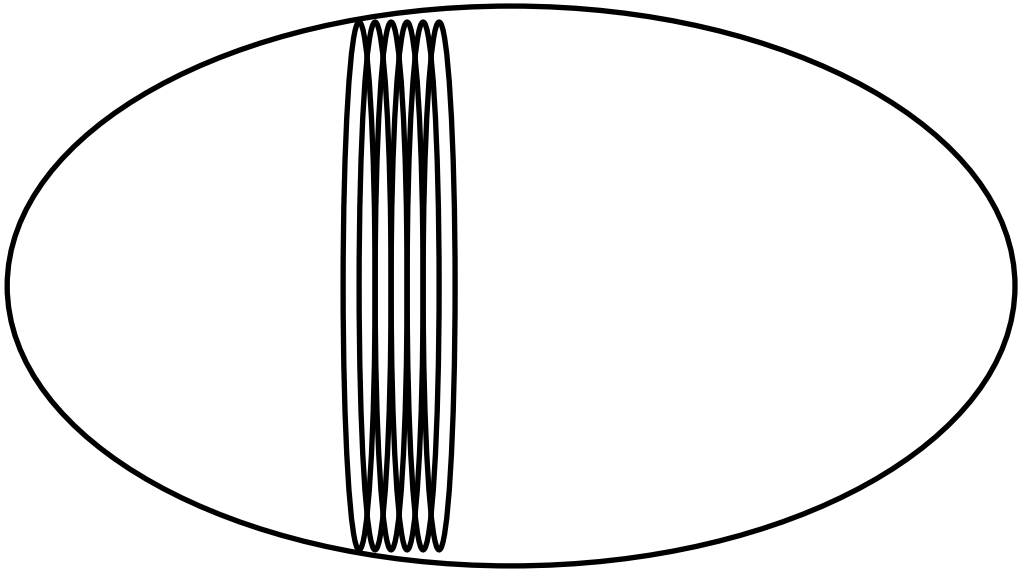

SWINBURNE UNIVERSITY OF TECHNOLOGY

Fold fan-beam mode

- Coherently dedisperse any pulsar in the beam.
- Up to ~30 pulsars in one beam
- Time > 500 pulsars/day

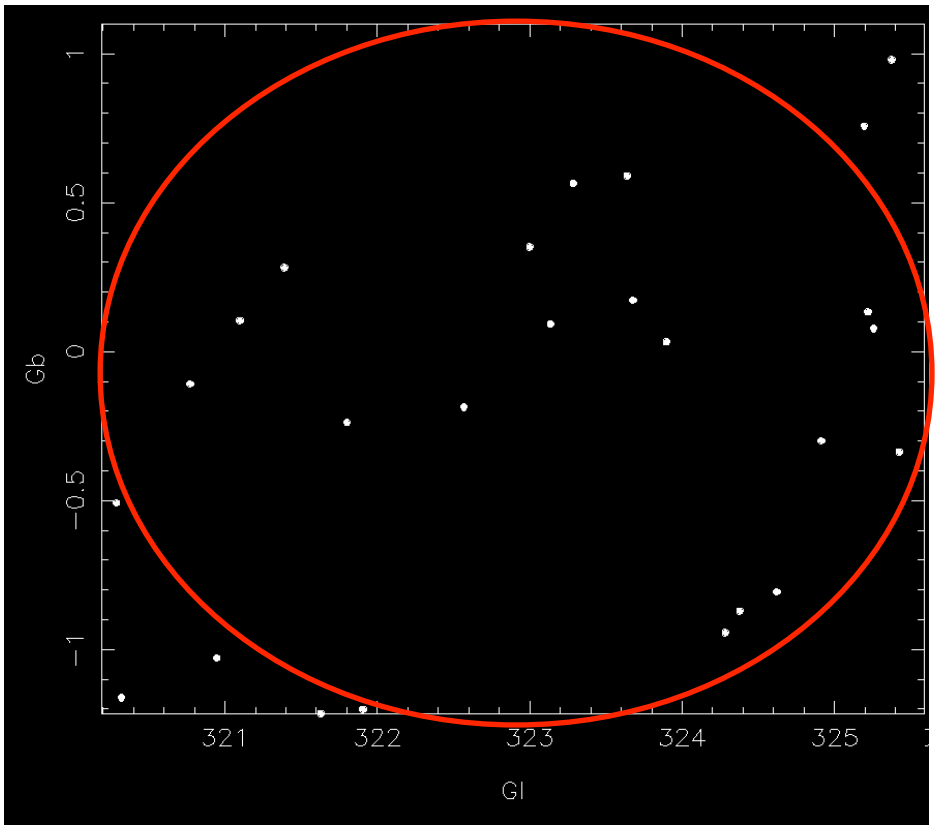

SWINBURNE UNIVERSITY OF TECHNOLOGY

New system vs old MOST system

- 4 x FoV
- 10 x Bandwidth
- 256 us timescale RFI monitor & excision
- 512 fan beams vs 96 or FX mode
- 15,000 spectral channels vs 1
- Multiple pulsar coherent dedispersion vs 1x3 MHz channel, one pulsar system
- *Simultaneous burst, mapping, pulsar mode, RFI excision modes using 33 dual-GPUs.*

SWINBURNE UNIVERSITY OF TECHNOLOGY

Timeline

- Correlator being debugged now!
- Incoherent modes working: ie 52 simultaneous Vela

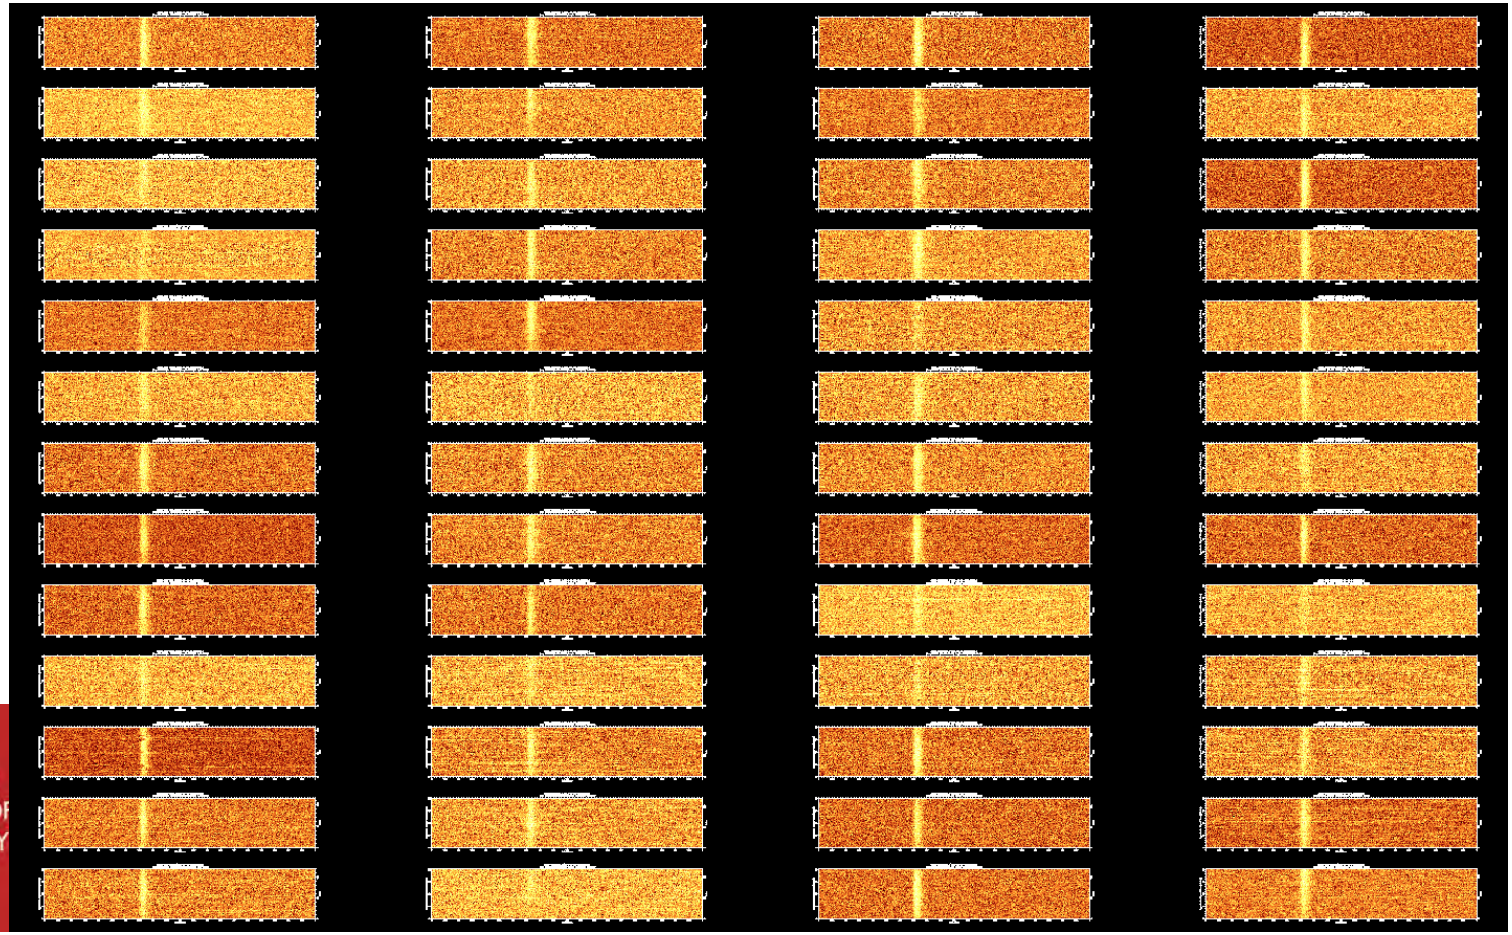

SWINBURNE UNIVERSITY OF TECHNOLOGY

Fringes

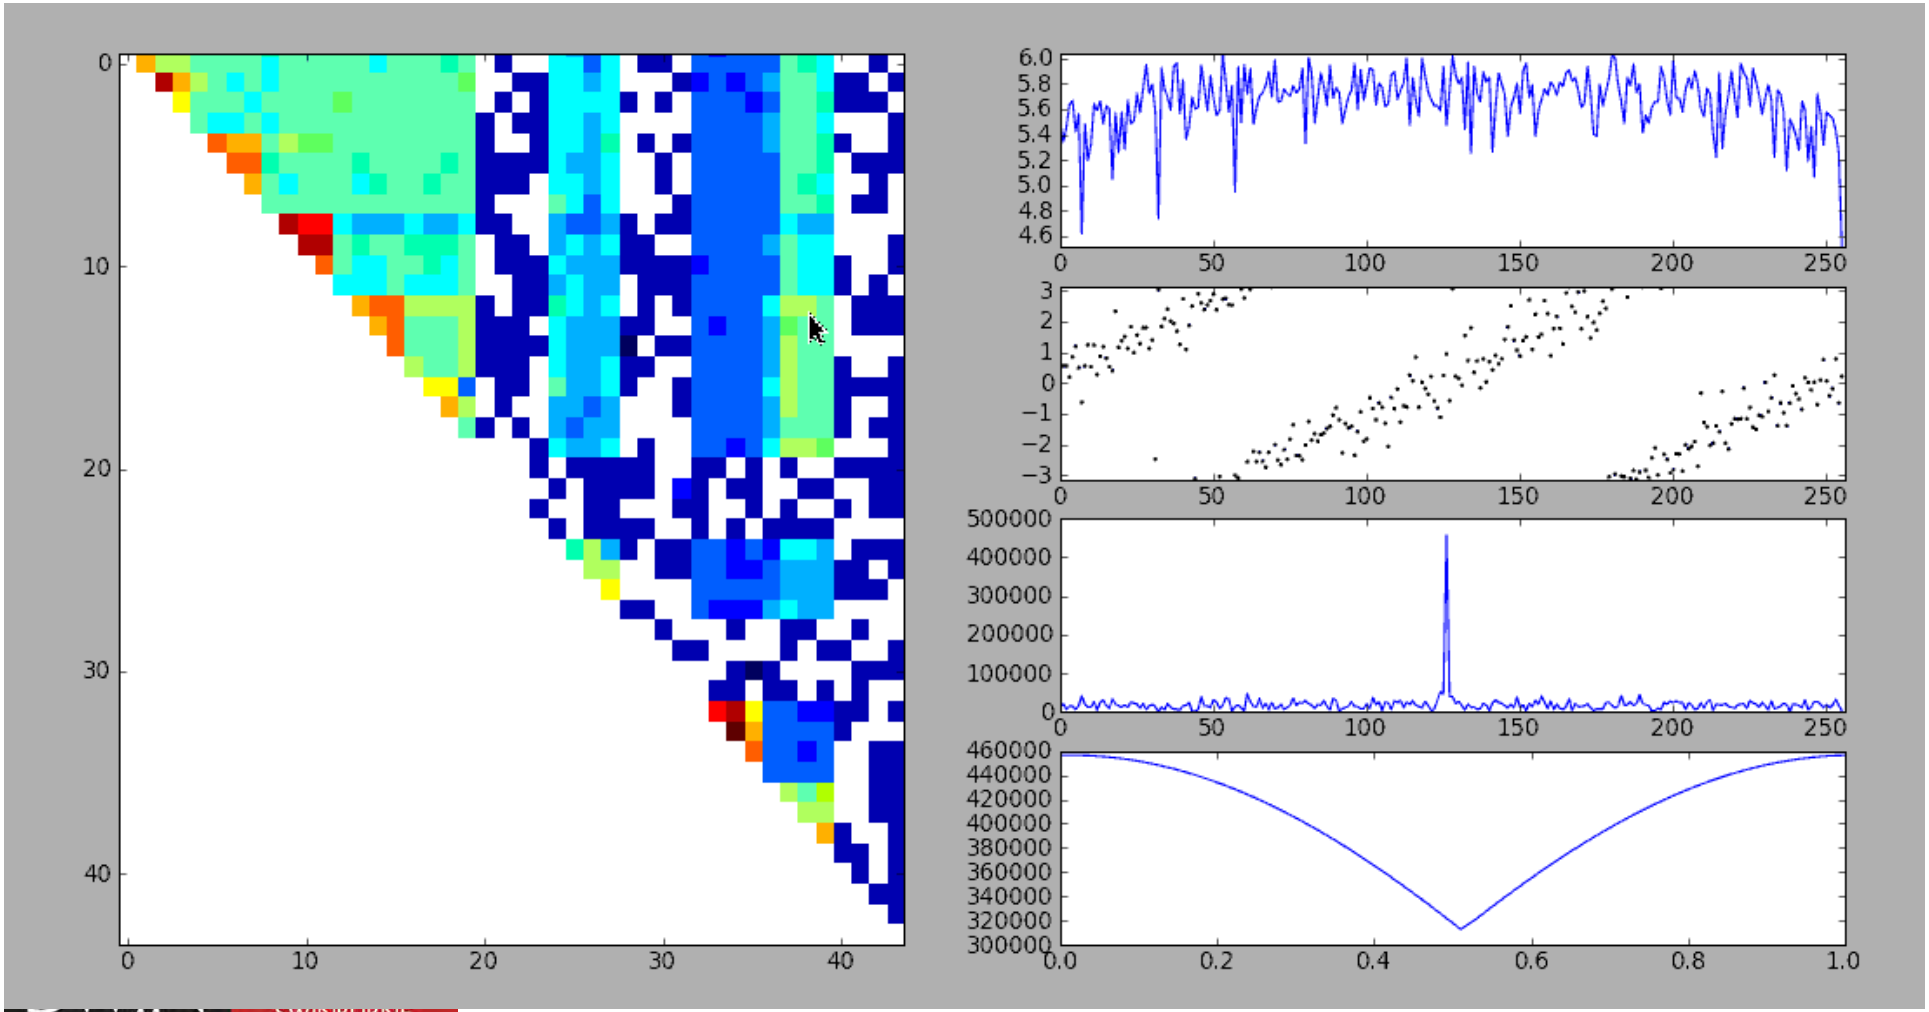

SWINBURNE UNIVERSITY OF TECHNOLOGY