Digital Backend Technology

Willem van Straten

Centre for Astrophysics & Supercomputing

SWINBURNE UNIVERSITY OF TECHNOLOGY

ATNF Radio Astronomy School 2009

What's not Digital?

- Mechanical Systems:
  - reflector, drive, …
- Analog Systems:
  - receiver, mixer, filter, chart recorder, …
- Organic Systems:
  - human cognition

SWINBURNE UNIVERSITY OF TECHNOLOGY


VA PROPERTY
ESTERLINE ANGUS

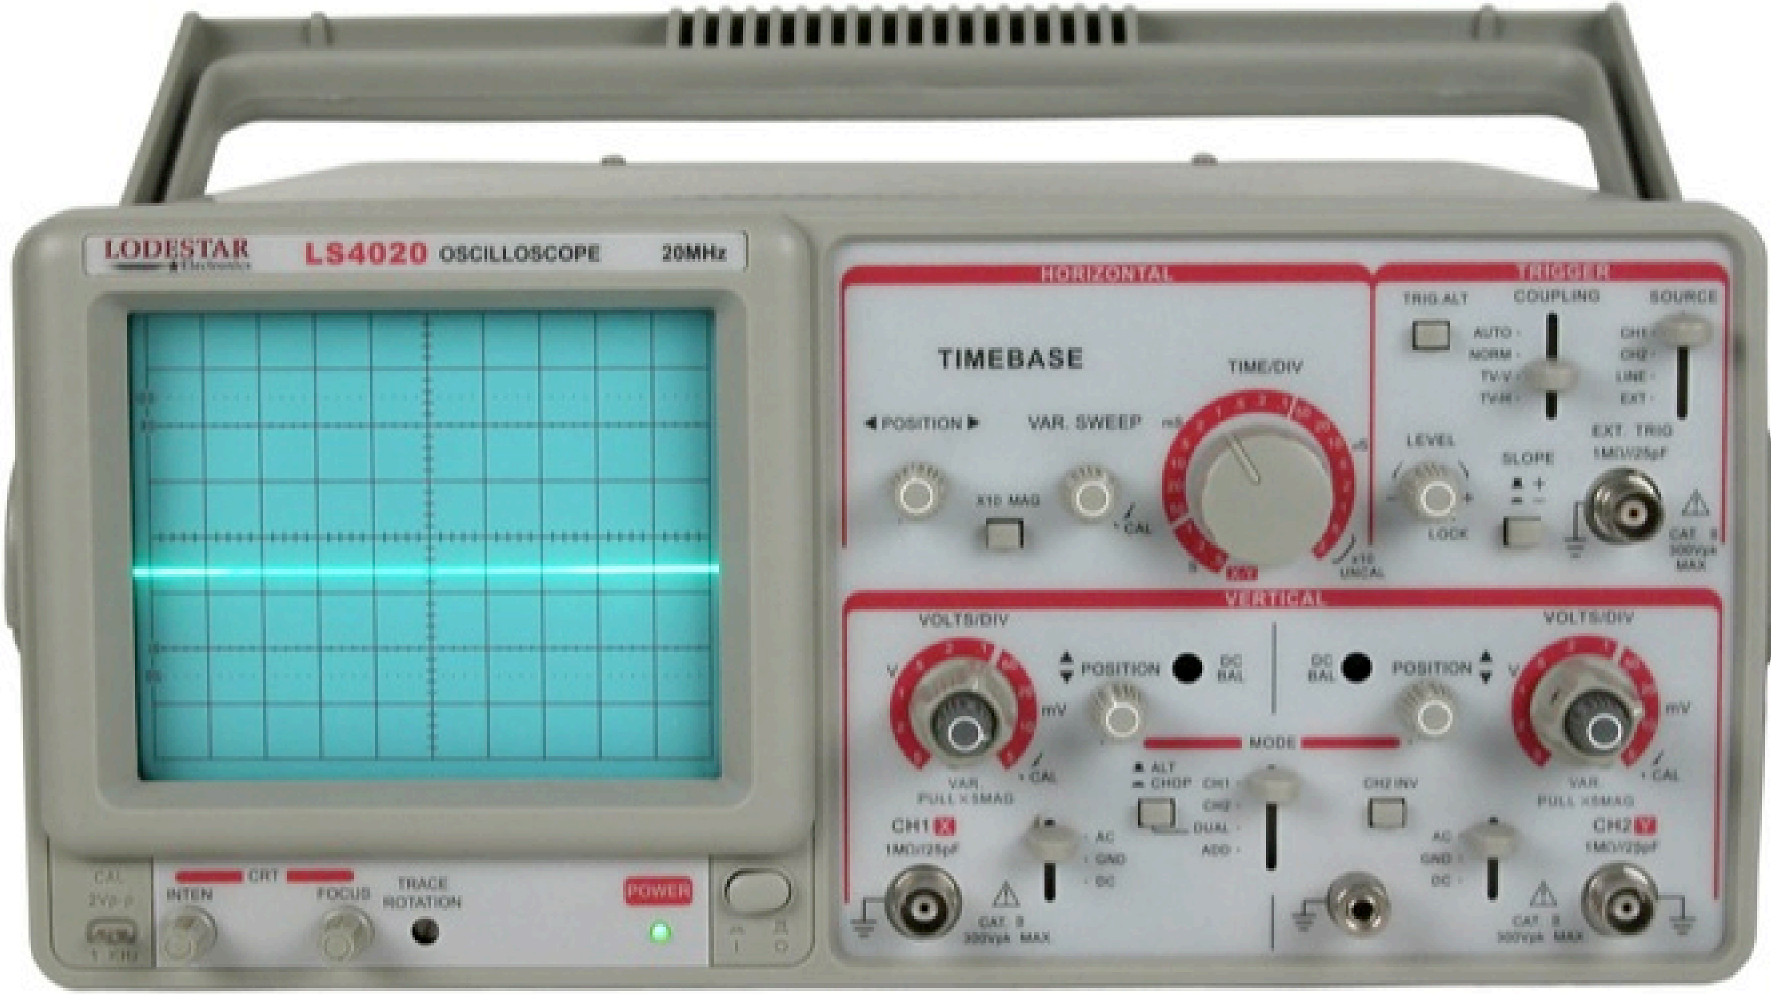
LODESTAR
LS4020 OSCILLOSCOPE
20MHz
HORIZONTAL
TRIGGER
TIMEBASE
TIME/DIV
POSITION
VAR. SWEEP
X10 MAG
CAL
UNCAL
TRIG.ALT
COUPLING
SOURCE
AUTO
NORM
TV-V
TV-H
CH1
CH2
LINE
EXT
LEVEL
LOCK
SLOPE
EXT. TRIG
VERTICAL
VOLTS/DIV
POSITION
DC BAL
MODE
ALT
CHOP
DUAL
ADD
CH2 INV
VAR
PULL ×5MAG
CH1
CH2
AC
GND
DC
CRT
INTEN
FOCUS
TRACE ROTATION
POWER

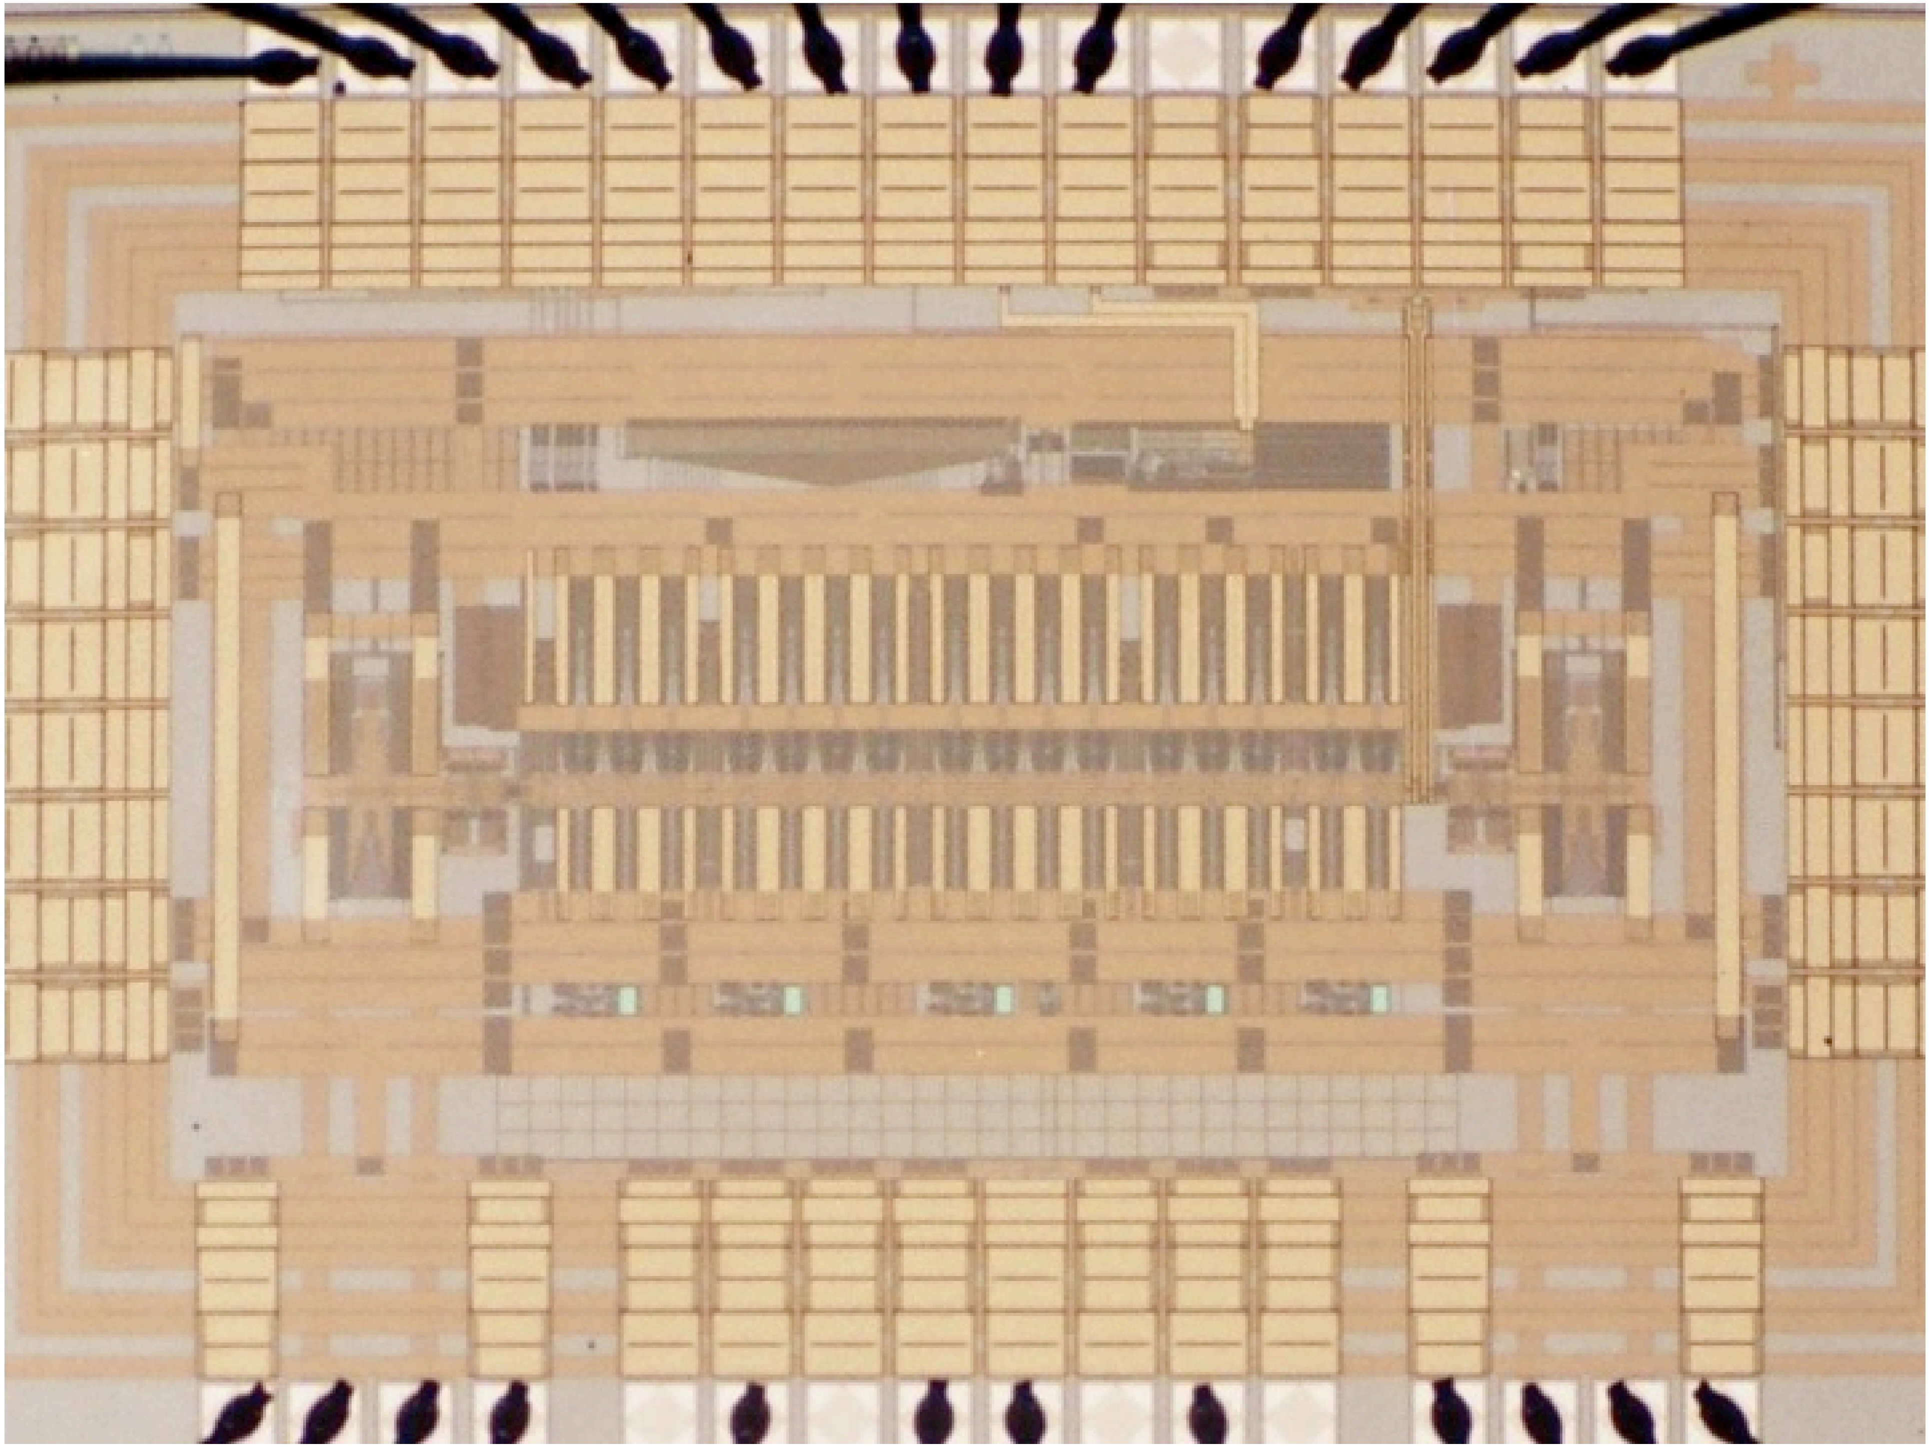

PROVE YOU'RE NOT A ROBOT

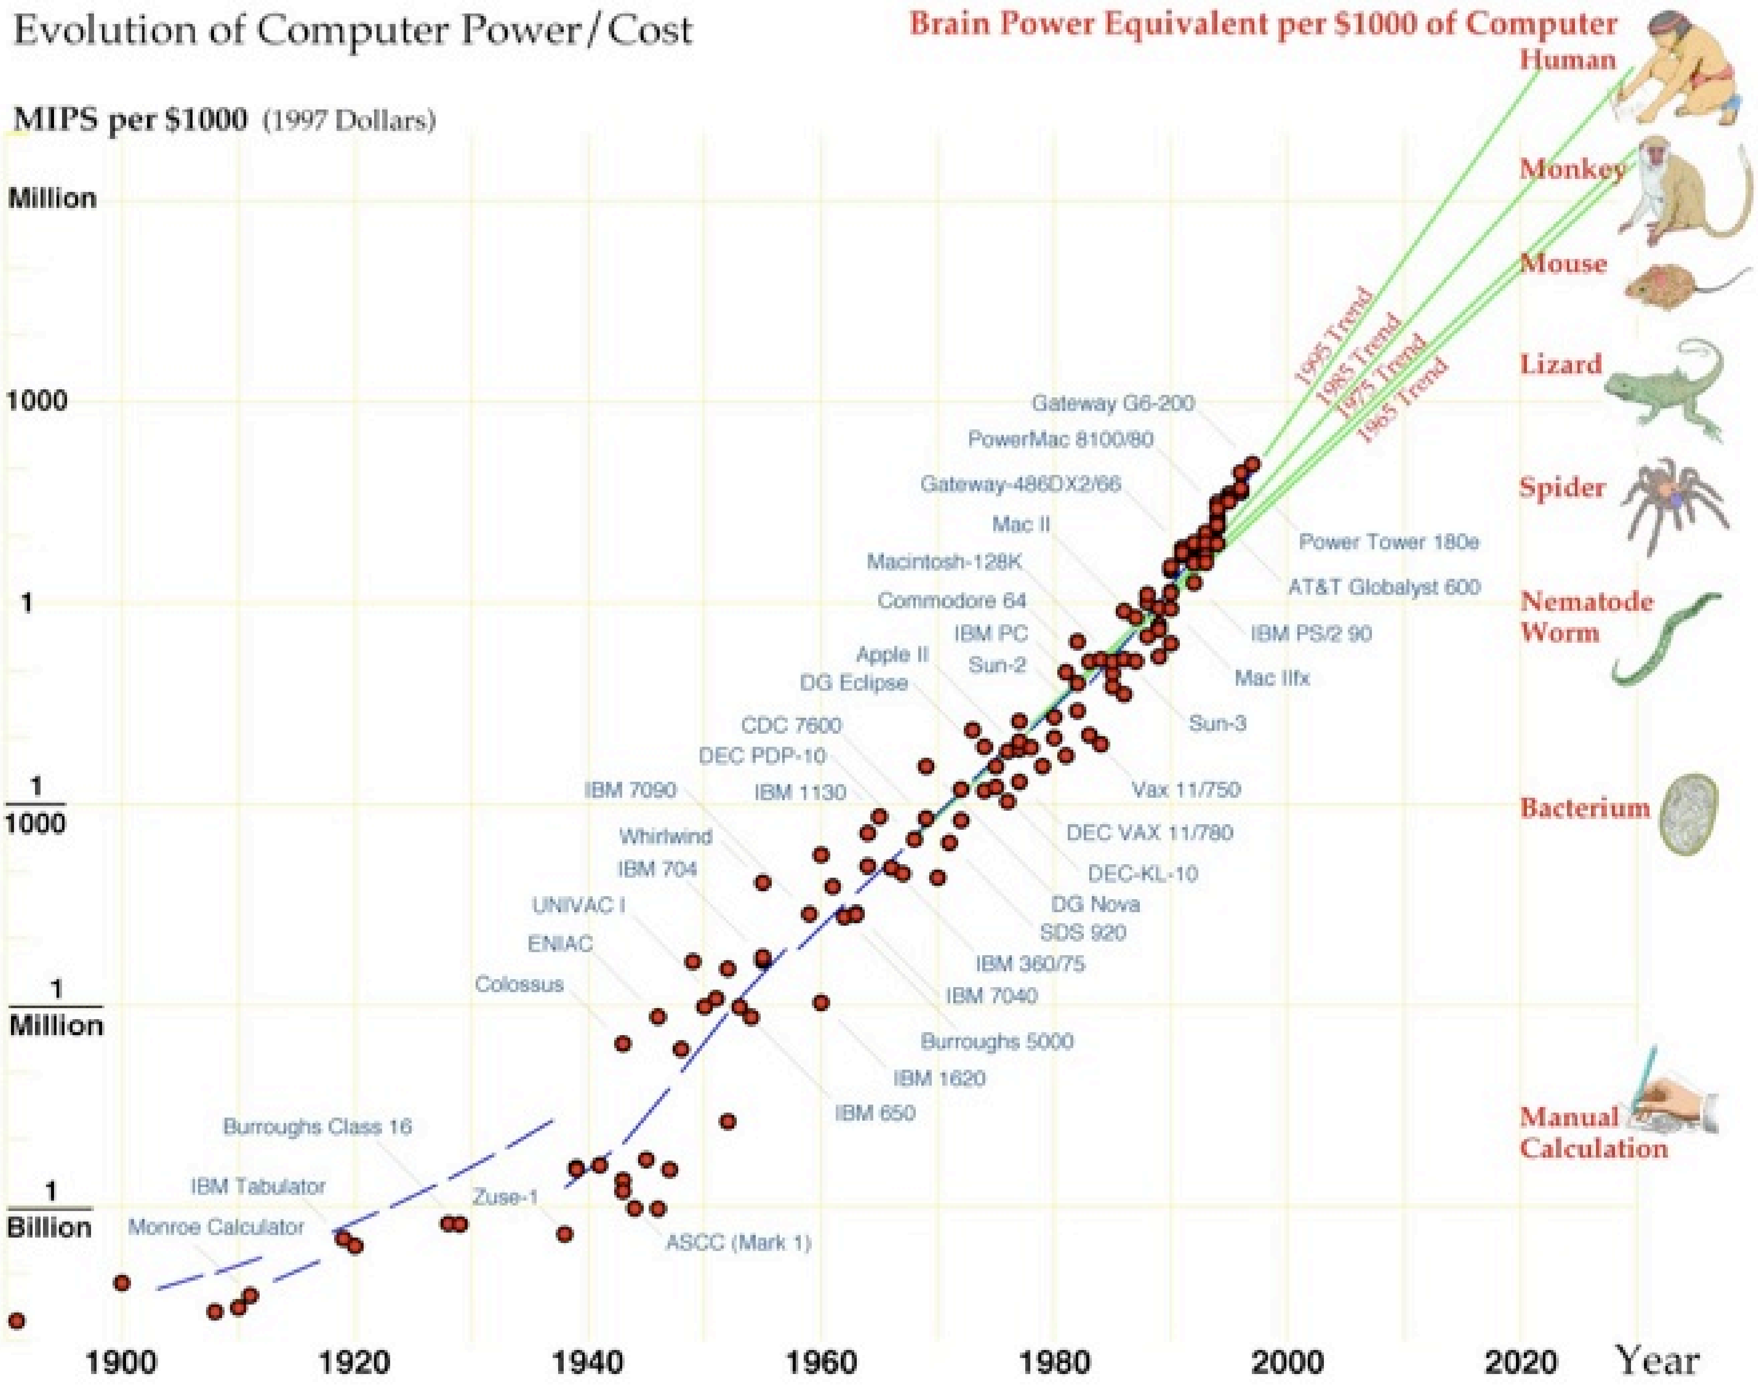
Evolution of Computer Power/Cost
MIPS per $1000 (1997 Dollars)
Brain Power Equivalent per $1000 of Computer
Million
1000
1
1/1000
1/Million
1/Billion
1900
1920
1940
1960
1980
2000
2020
Year
Human
Monkey
Mouse
Lizard
Spider
Nematode Worm
Bacterium
Manual Calculation
1995 Trend
1985 Trend
1975 Trend
1965 Trend
Gateway G6-200
PowerMac 8100/80
Gateway-486DX2/66
Mac II
Macintosh-128K
Commodore 64
IBM PC
Apple II
Sun-2
DG Eclipse
CDC 7600
DEC PDP-10
IBM 7090
IBM 1130
Whirlwind
IBM 704
UNIVAC I
ENIAC
Colossus
Burroughs Class 16
IBM Tabulator
Zuse-1
Monroe Calculator
ASCC (Mark 1)
Power Tower 180e
AT&T Globalyst 600
IBM PS/2 90
Mac IIfx
Sun-3
Vax 11/750
DEC VAX 11/780
DEC-KL-10
DG Nova
SDS 920
IBM 360/75
IBM 7040
Burroughs 5000
IBM 1620
IBM 650

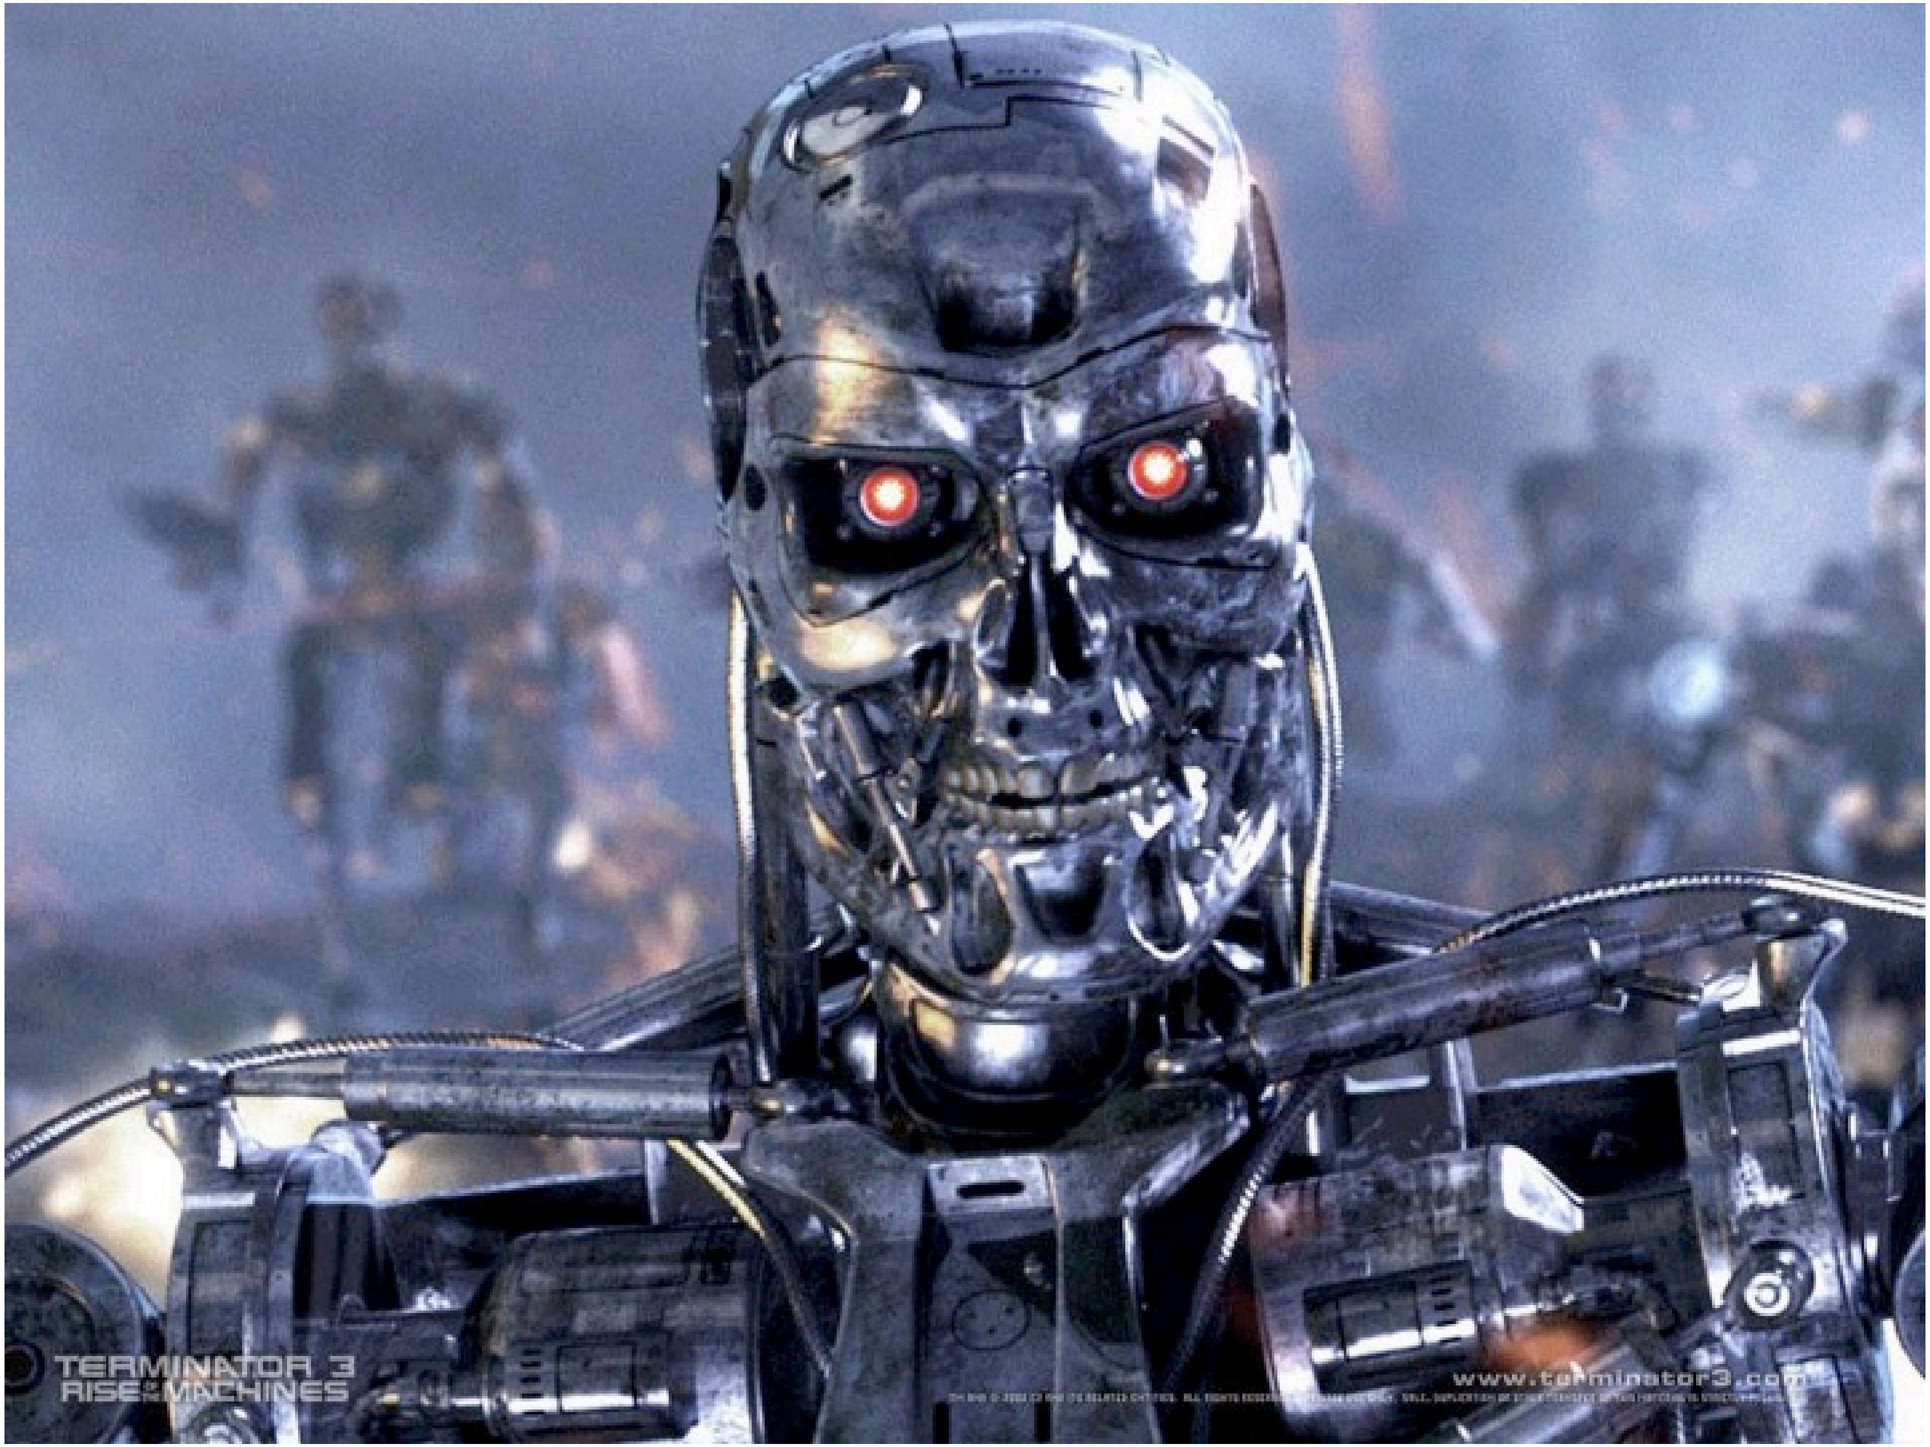
TERMINATOR 3
RISE MACHINES
www.terminator3.com

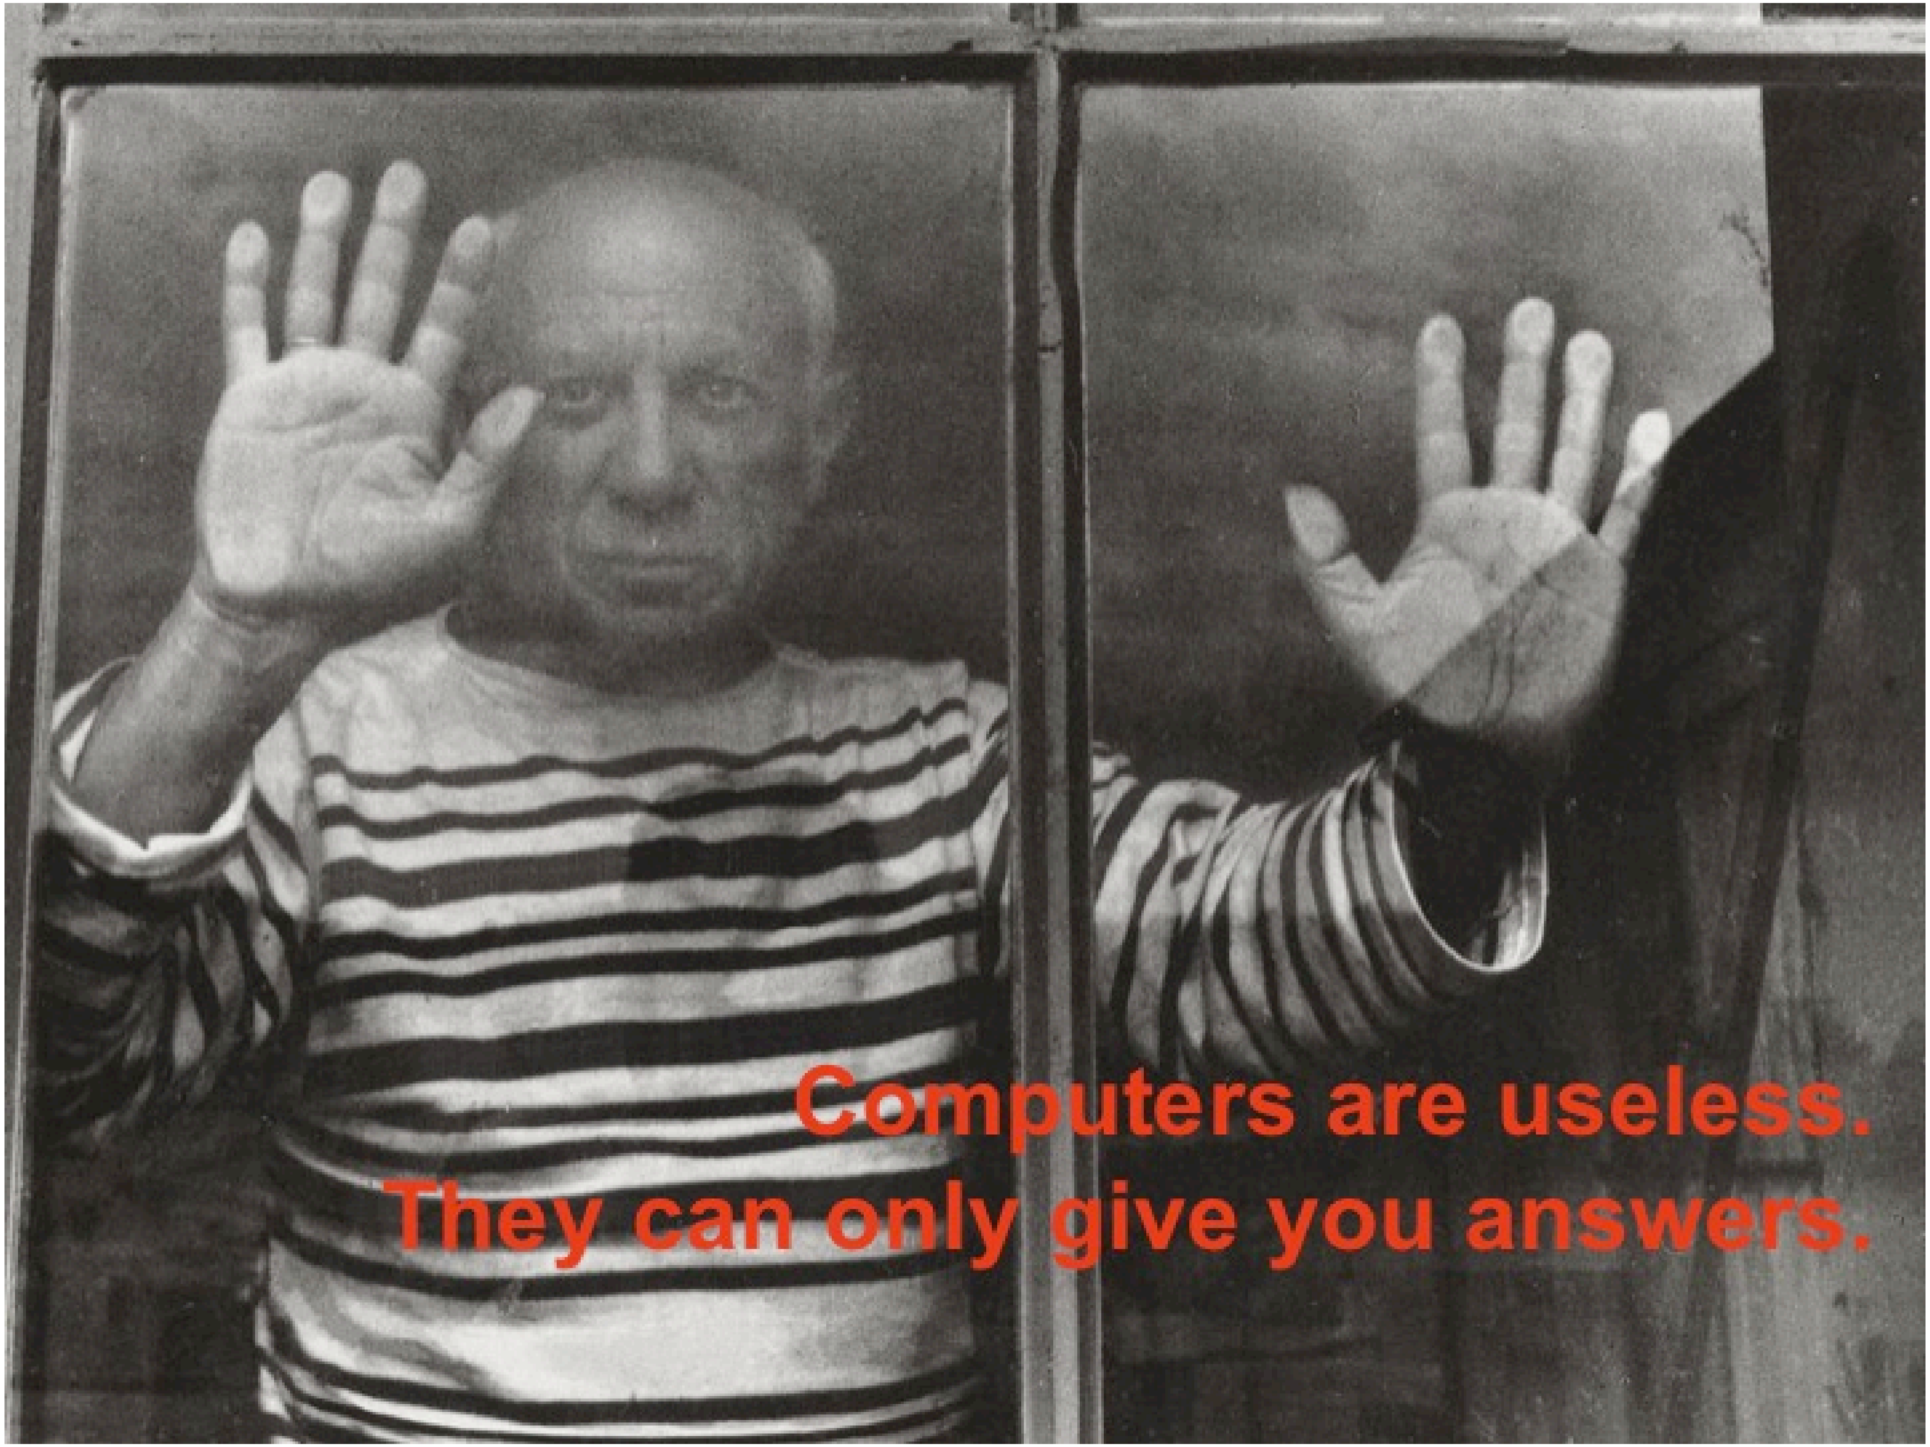
Computers are useless.
They can only give you answers.

Digital Data Reduction

- Common constraints:
  - processing speed, bandwidth, capacity, cost
- Different technologies:
  - ASIC, FPGA, CPU, GPU, etc.
- The current state of the art
  - synergy between hardware/software
- The Future

SWINBURNE UNIVERSITY OF TECHNOLOGY

Common Constraints - Nyquist

- Analog-to-digital conversion

  Nyquist theorem: sampling rate = 2 * bandwidth

- Many astronomical signals are broadband

  (decametre to millimetre wavelengths)

- Signal-to-noise ratio (S/N)

  proportional to sqrt (time * bandwidth)

Common Constraints - Dynamic Range

- Data rate = 2 * BW * npol * nbit   [bits/s]

- Early systems used 2 bits / sample
  - significant quantization error (noise and distortion)

- Modern need for dynamic range
  - mostly due to radio frequency interference (RFI)
  - 8 bits / sample now common

SWINBURNE UNIVERSITY OF TECHNOLOGY

Common Constraints - Data Transport

- Transfer to memory of computational device
- Direct memory access (e.g. via PCI or VME bus)
- Ethernet (e.g. TCP/IP or UDP/IP)
- Intermediate storage (e.g. magnetic tape or disk)
- Memory bandwidth (b/w memory and processor)

Common Constraints - Processing

- Processor speed (MIPS, FLOPS)

- Real time vs offline:
  - Data storage facility (speed, capacity)
  - duty cycle: acquisition vs processing time
  - local or off-site computing resources

- Storage/handling of results:
  - Off site, tape archive

SWINBURNE UNIVERSITY OF TECHNOLOGY

Common Constraints - Cost

- Initial technology purchase
- Research and development:
  - personnel and infrastructure
- Operational costs:
  - power consumption, cooling requirements, space
  - system administration and maintenance
  - data dissemination (web server)

Digital Backend Technology

- Hardware:
  - increasing flexibility and modularity
- Software:
  - increasing performance
- Distinction becoming increasingly blurry

SWINBURNE UNIVERSITY OF TECHNOLOGY

Hardware-based Technologies

- Digital Signal Processor (DSP)
  - microprocessor optimized for digital signal processing
  - low latency, direct memory access, fast multiply-accumulate, etc.
- Application-specific Integrated Circuit (ASIC)
  - Hardware Description Language (HDL) design tools
  - design sent to manufacturer
- Field-programmable Gate Array (FPGA)
  - programmable logic (useful during R&D)
  - more flexible than ASICs

Software-based Technologies

- Central Processing Unit (CPU) [~Gflops]
  - execute computer programs (instruction sets)
  - high-level languages such as C++, Java, python, Matlab, etc.
- Graphics Processing Unit (GPU) [~Tflops]
  - Hardware graphics accelerators
  - CUDA, OpenCL
- Parallel computing [~Pflops]
  - multi-core, multi-processor computers (e.g. threads)
  - clusters, grids (e.g. Message Passing Interface)

What do backends do?

- Filter:
  - convolution by finite impulse response (FIR)
- Spectrometry:
  - ACF, FFT, polyphase filterbank, etc.
- Integration of statistical quantities:
  - for pulsar work, phase-resolved average

Two Examples at Parkes

- Pulsar Instrumentation:
  - need for time and frequency resolution
  - synthesis of hardware and software

- High-precision timing
  - real-time data reduction

- Pulsar survey
  - offline analysis on super computer

SWINBURNE UNIVERSITY OF TECHNOLOGY

Radio Pulsars

- Rapidly spinning, highly magnetized neutron stars:
  - M ~ 1.4 solar masses
  - D ~ 20 km
  - B ~ 10^8 - 10^{14} G
  - P ~ 1 ms – 10 s

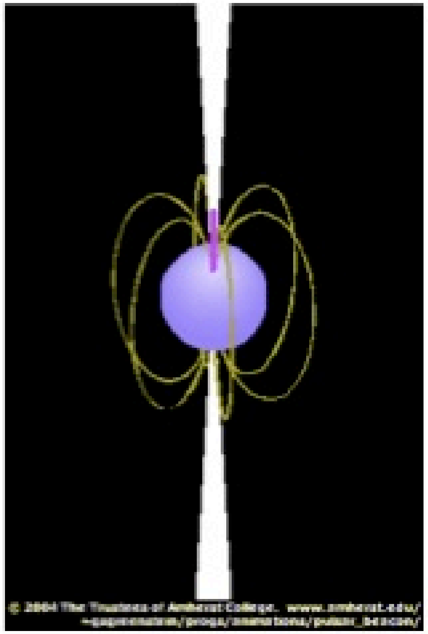

- Radio beams from magnetic poles
  - stellar lighthouse

SWIN BUR NE SWINBURNE UNIVERSITY OF TECHNOLOGY

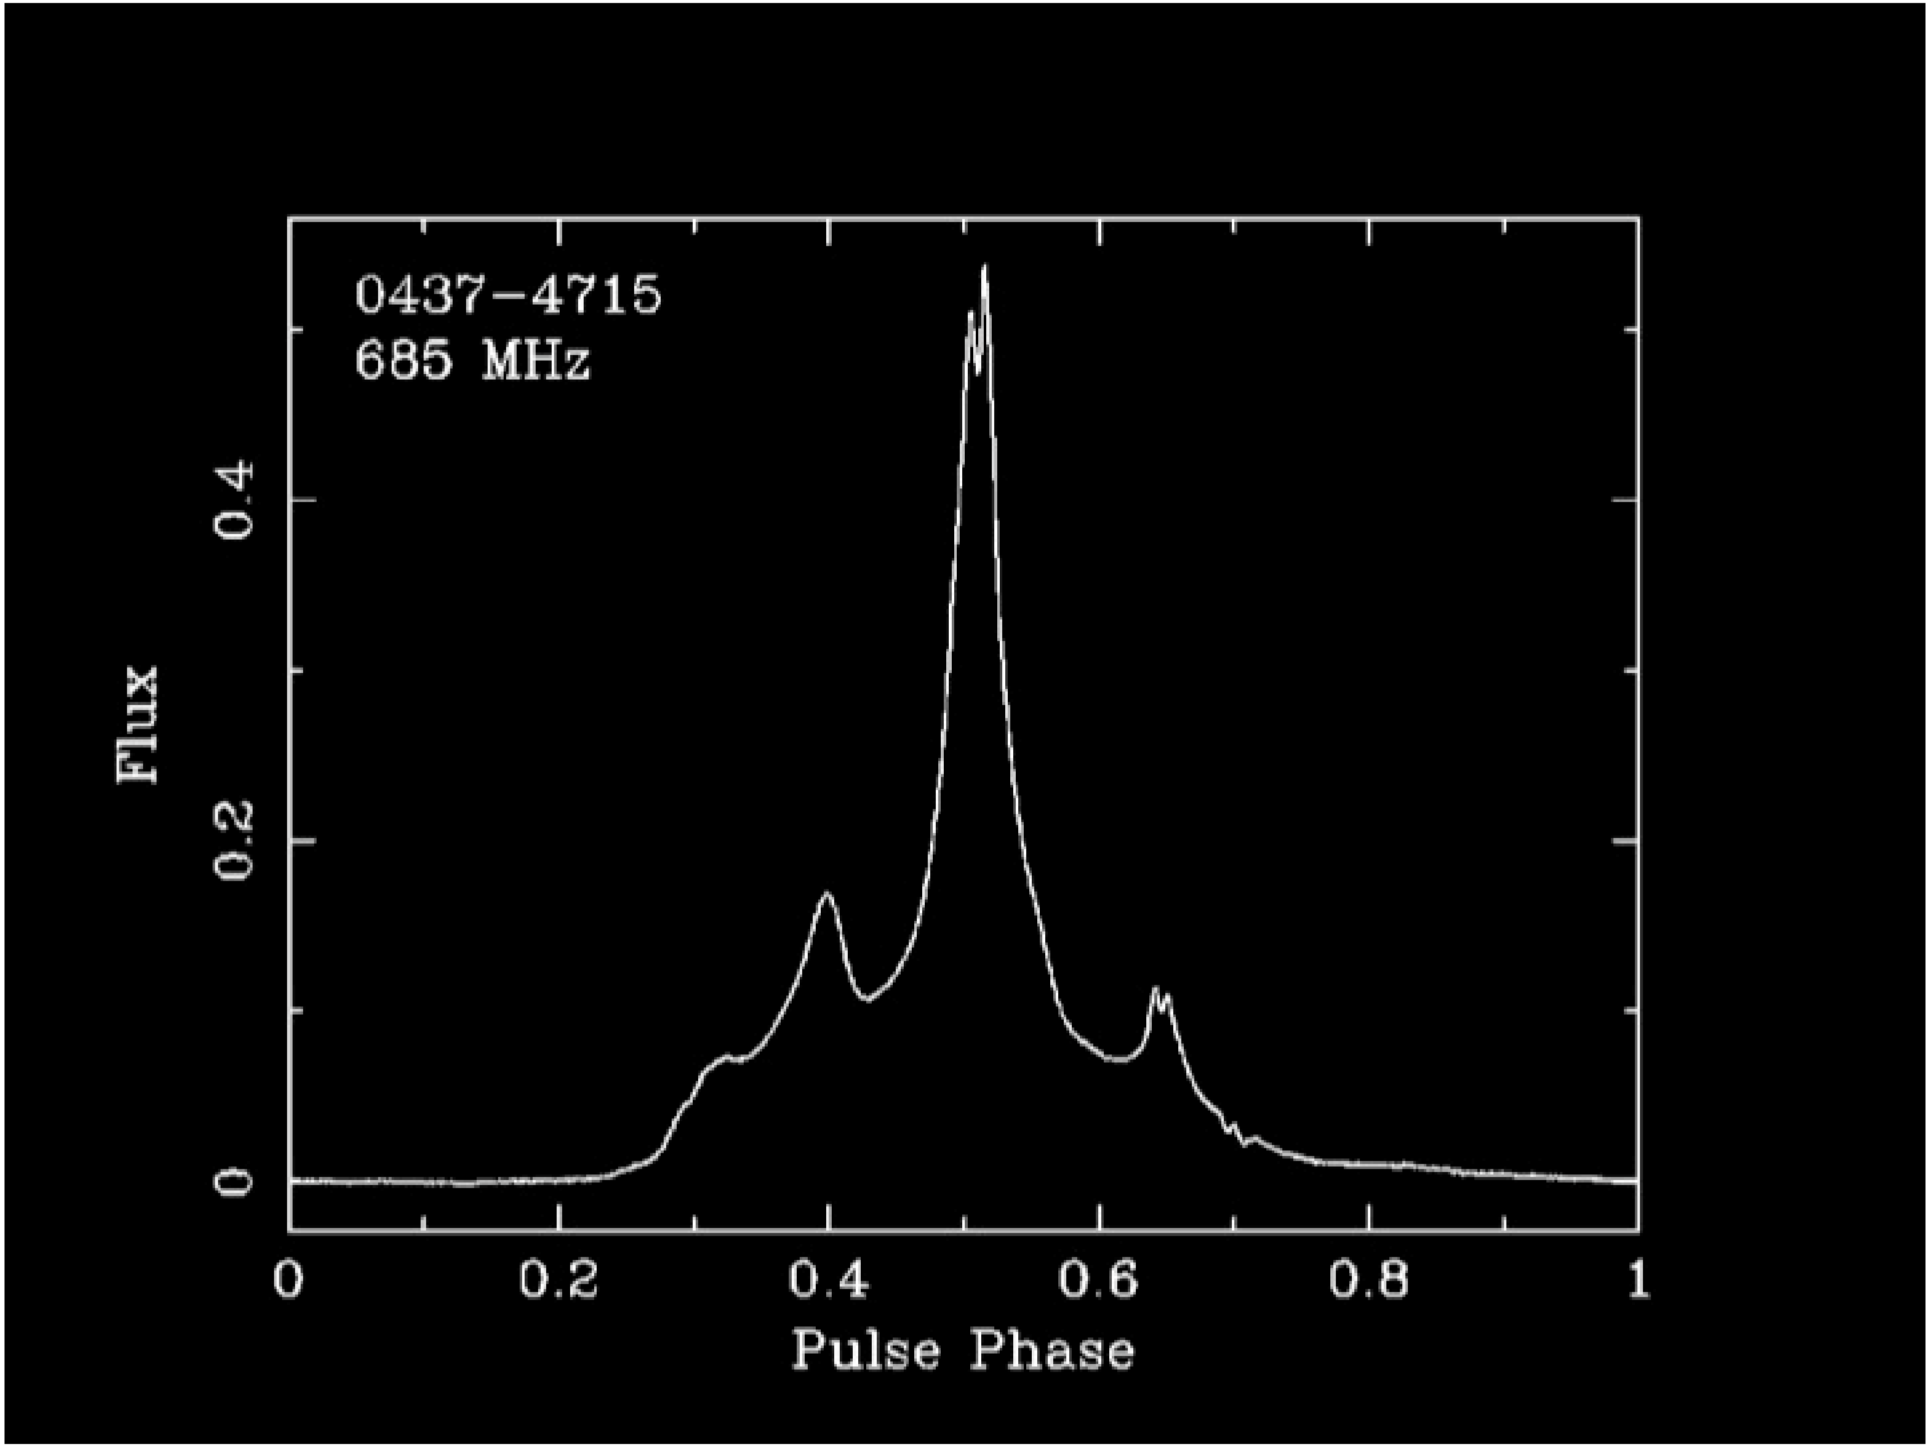
0437-4715
685 MHz
Flux
0
0.2
0.4
0
0.2
0.4
0.6
0.8
1
Pulse Phase

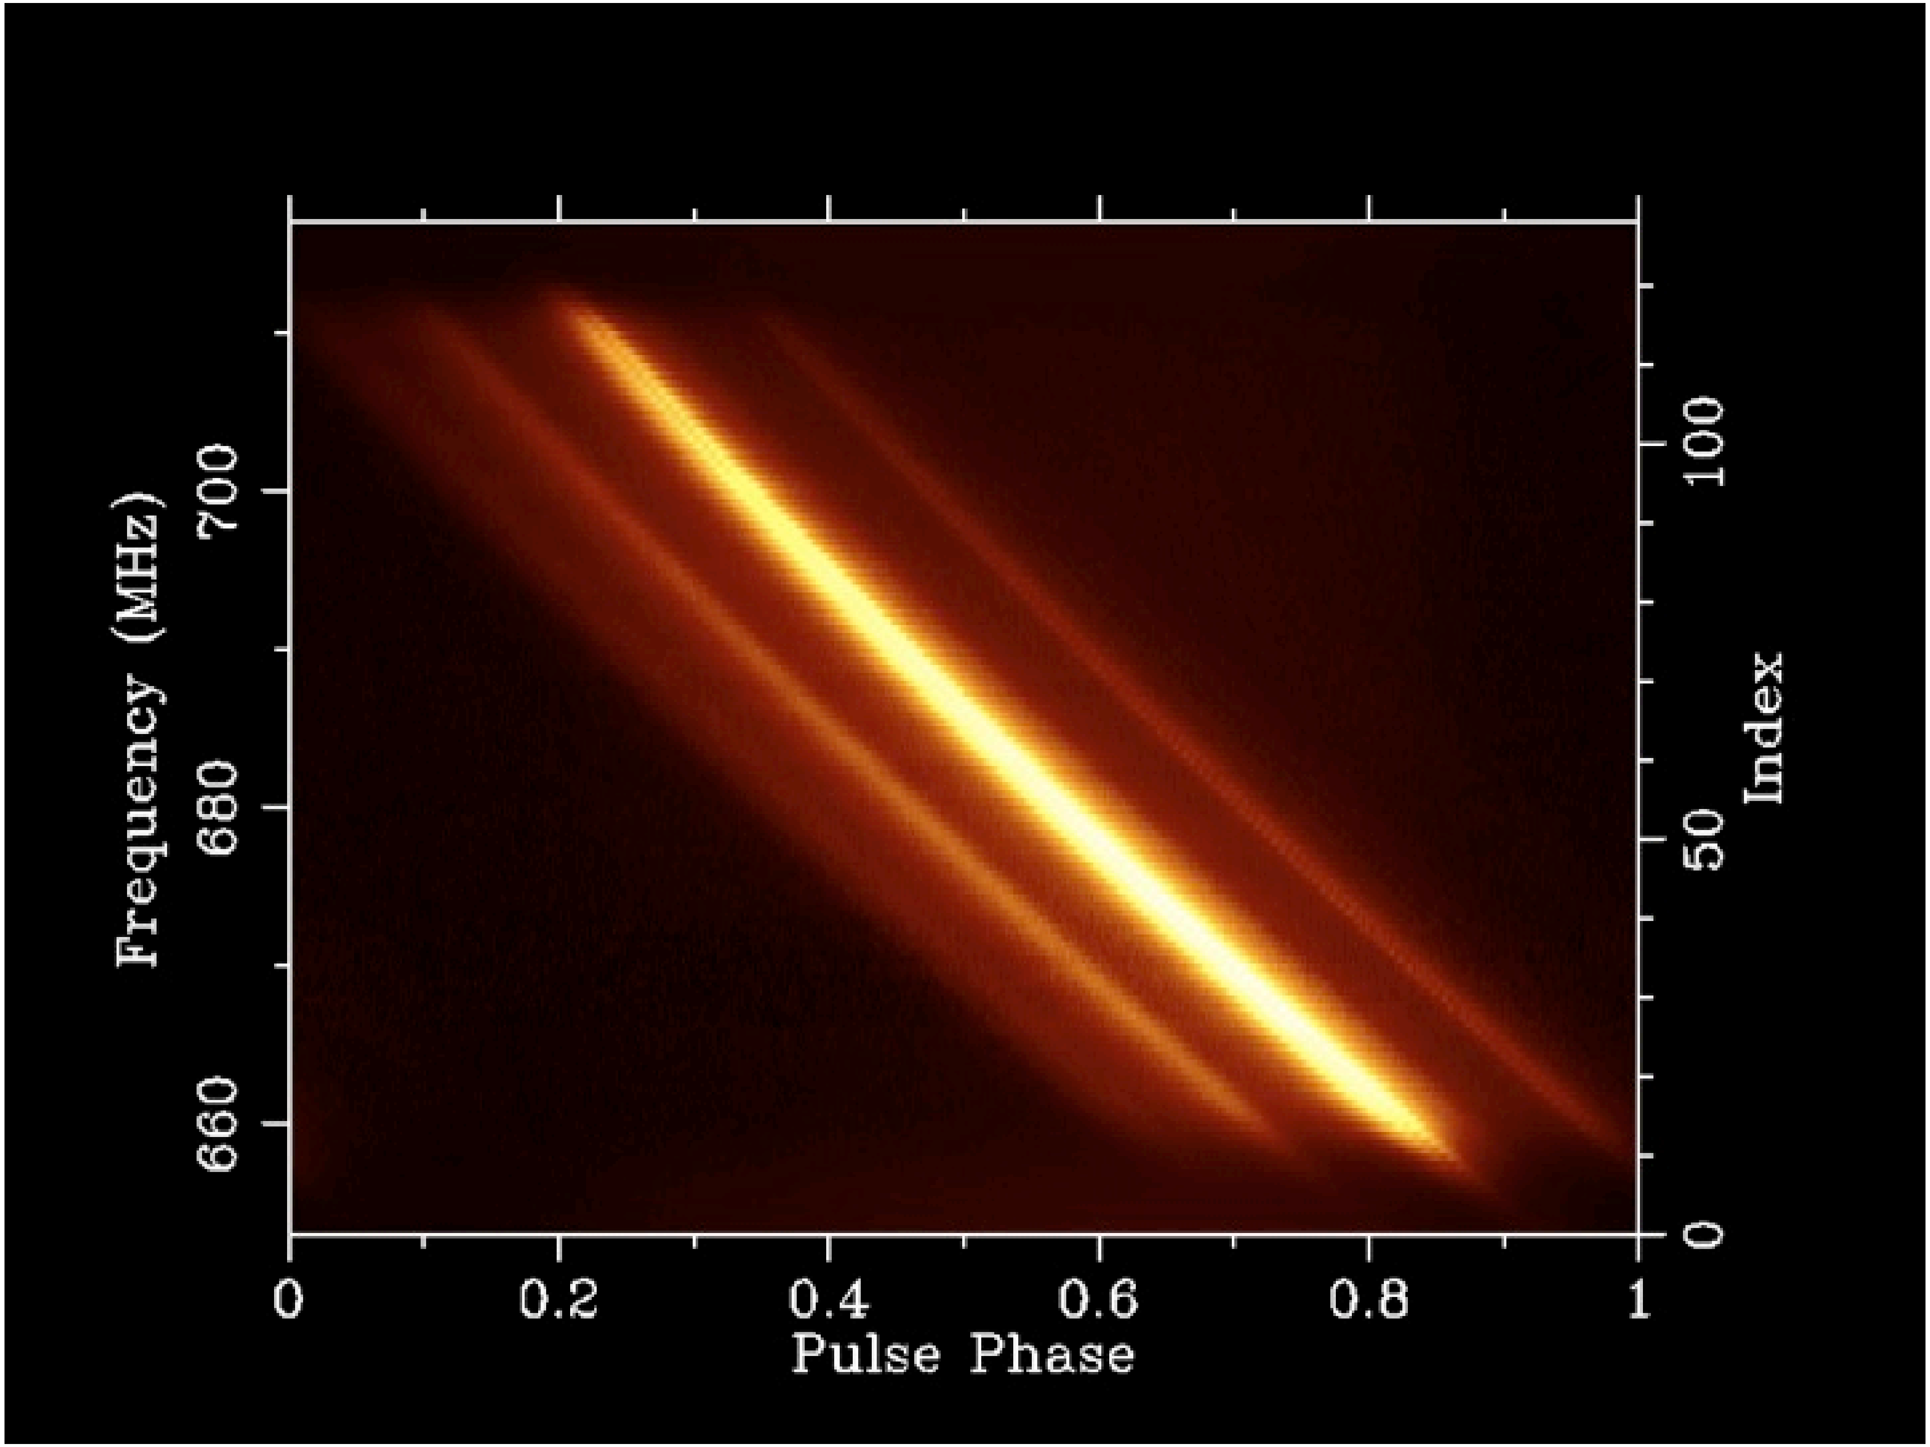

Frequency (MHz)
700
680
660
Index
100
50
0
0
0.2
0.4
0.6
0.8
1
Pulse Phase

Dispersion removal backends

- Incoherent (post-detection):
  - analog filterbanks
  - digital filterbanks:
    - autocorrelation spectrometers
    - Fast Fourier Transform
    - polyphase filterbank (e.g. MPEG audio)

- Phase-coherent (pre-detection):
  - baseband recording and processing systems
    - tape, disk, or real-time

Phase-coherent dispersion removal

- Observed voltage signal is deconvolved
  - impulse response function of ISM plasma dispersion

- Convolution performed in frequency domain
  - more efficient, requires FFT

- N_{fft} proportional to DM (and ~ $\delta\nu/\nu^3$)

Phase-coherent dispersion removal

- $O_{FFT}(N) = 5N\log N$ (FFT benchmark)
- flops = $5N\log N / t_{FFT}$
- DM=10 and $\delta\nu$ = 16 MHz
  - ν=1400MHz $\Rightarrow$ N=128k [1.5 Gflops]
  - ν=400MHz $\Rightarrow$ N=2M [1.8 Gflops]
- DM=10 and $\delta\nu$ = 128 MHz
  - ν=1400MHz $\Rightarrow$ N=8M [15.8 Gflops]
  - ν=400MHz $\Rightarrow$ N=256M [20 Gflops]

SWINBURNE UNIVERSITY OF TECHNOLOGY

128 MHz

16 MHz

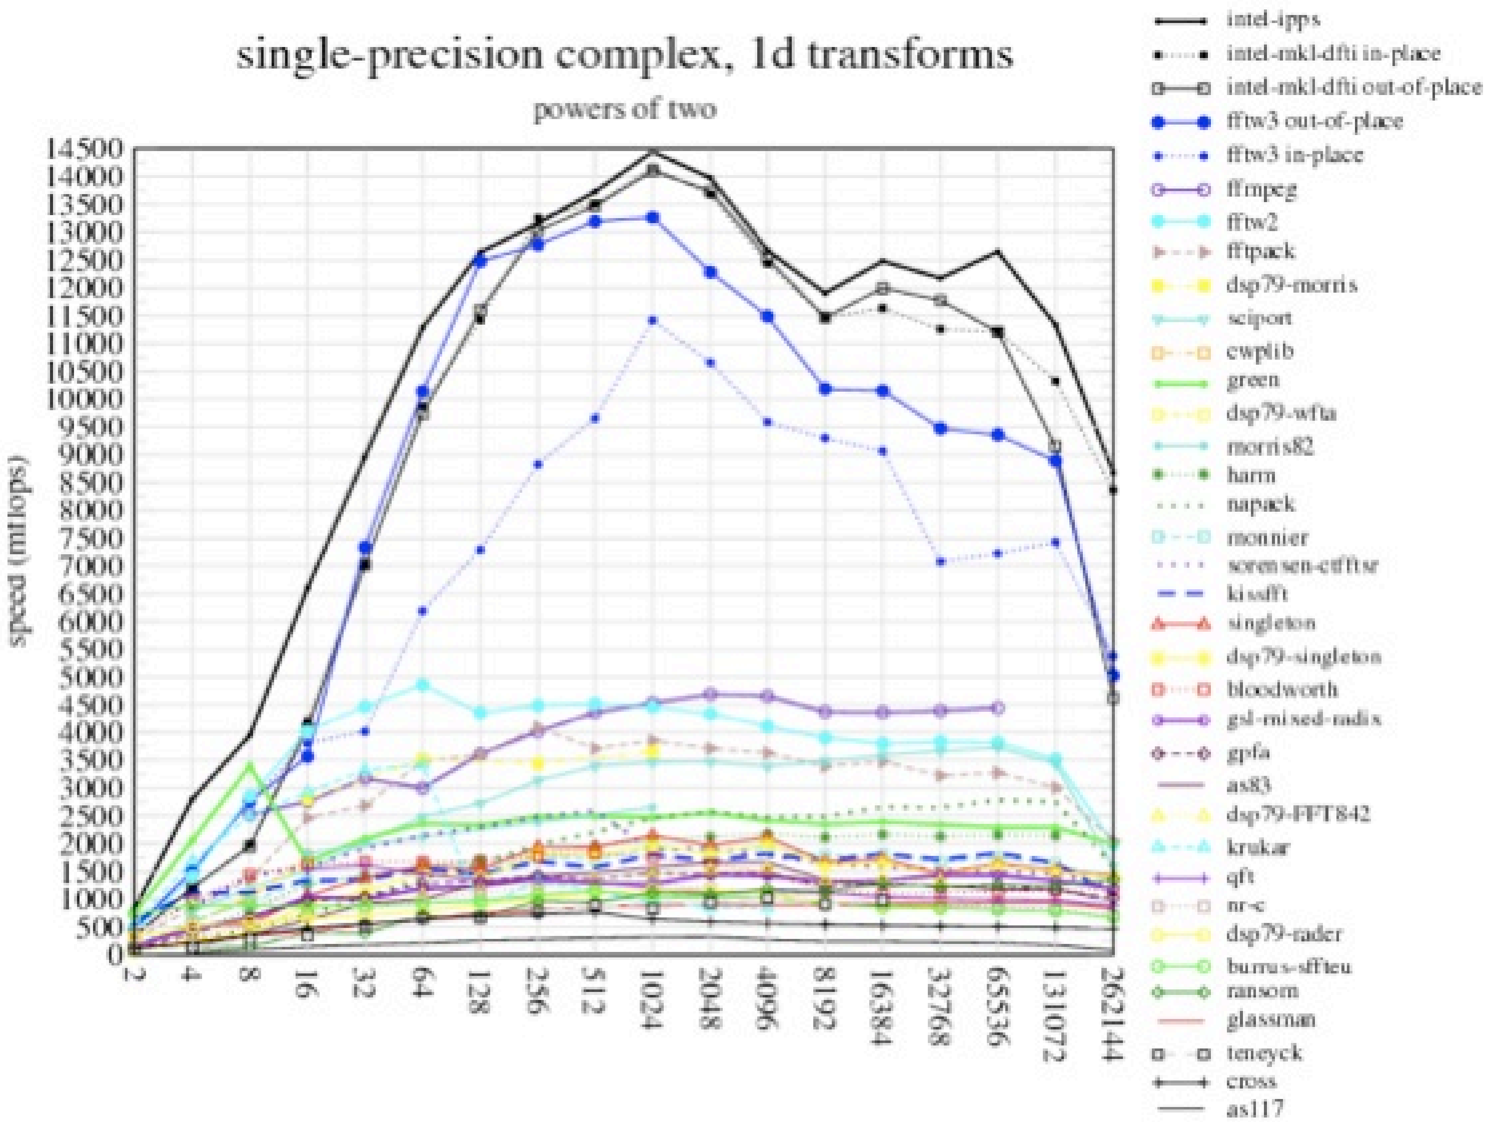

http://www.fftw.org/speed/CoreDuo-3.0GHz-icc/

Coherent Dedispersion - History

- 1971 - **0.125 MHz** @ Arecibo
  - XDS Sigma 5 magnetic tape
  - 20% duty cycle for 3 minutes
- 1987 - **1.5 MHz**
  - Reticon R5601 chip
  - real-time!
- 1998 - **16 MHz** @ Parkes
  - S2, VHS tape
- 1999 - **20 MHz** @ Parkes
  - CPSR, DLT tape
- 2002 - **128 MHz**
  - CPSR2, high-speed disk
- 2007 - **1024 MHz**
  - APSR, real-time

SWINBURNE UNIVERSITY OF TECHNOLOGY

Coherent Dedispersion - Friday

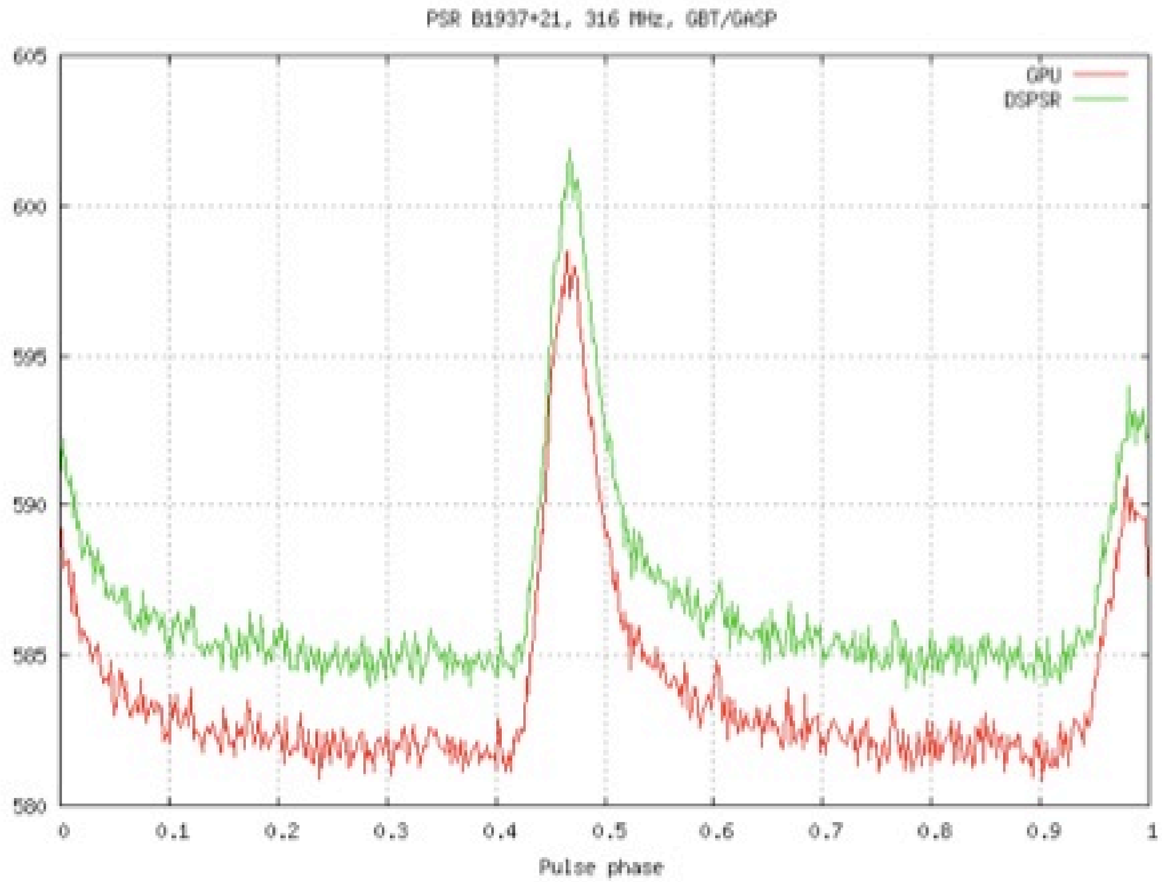

100 MHz on one GPU

SWINBURNE UNIVERSITY OF TECHNOLOGY


XDS SIGMA 5
CPU

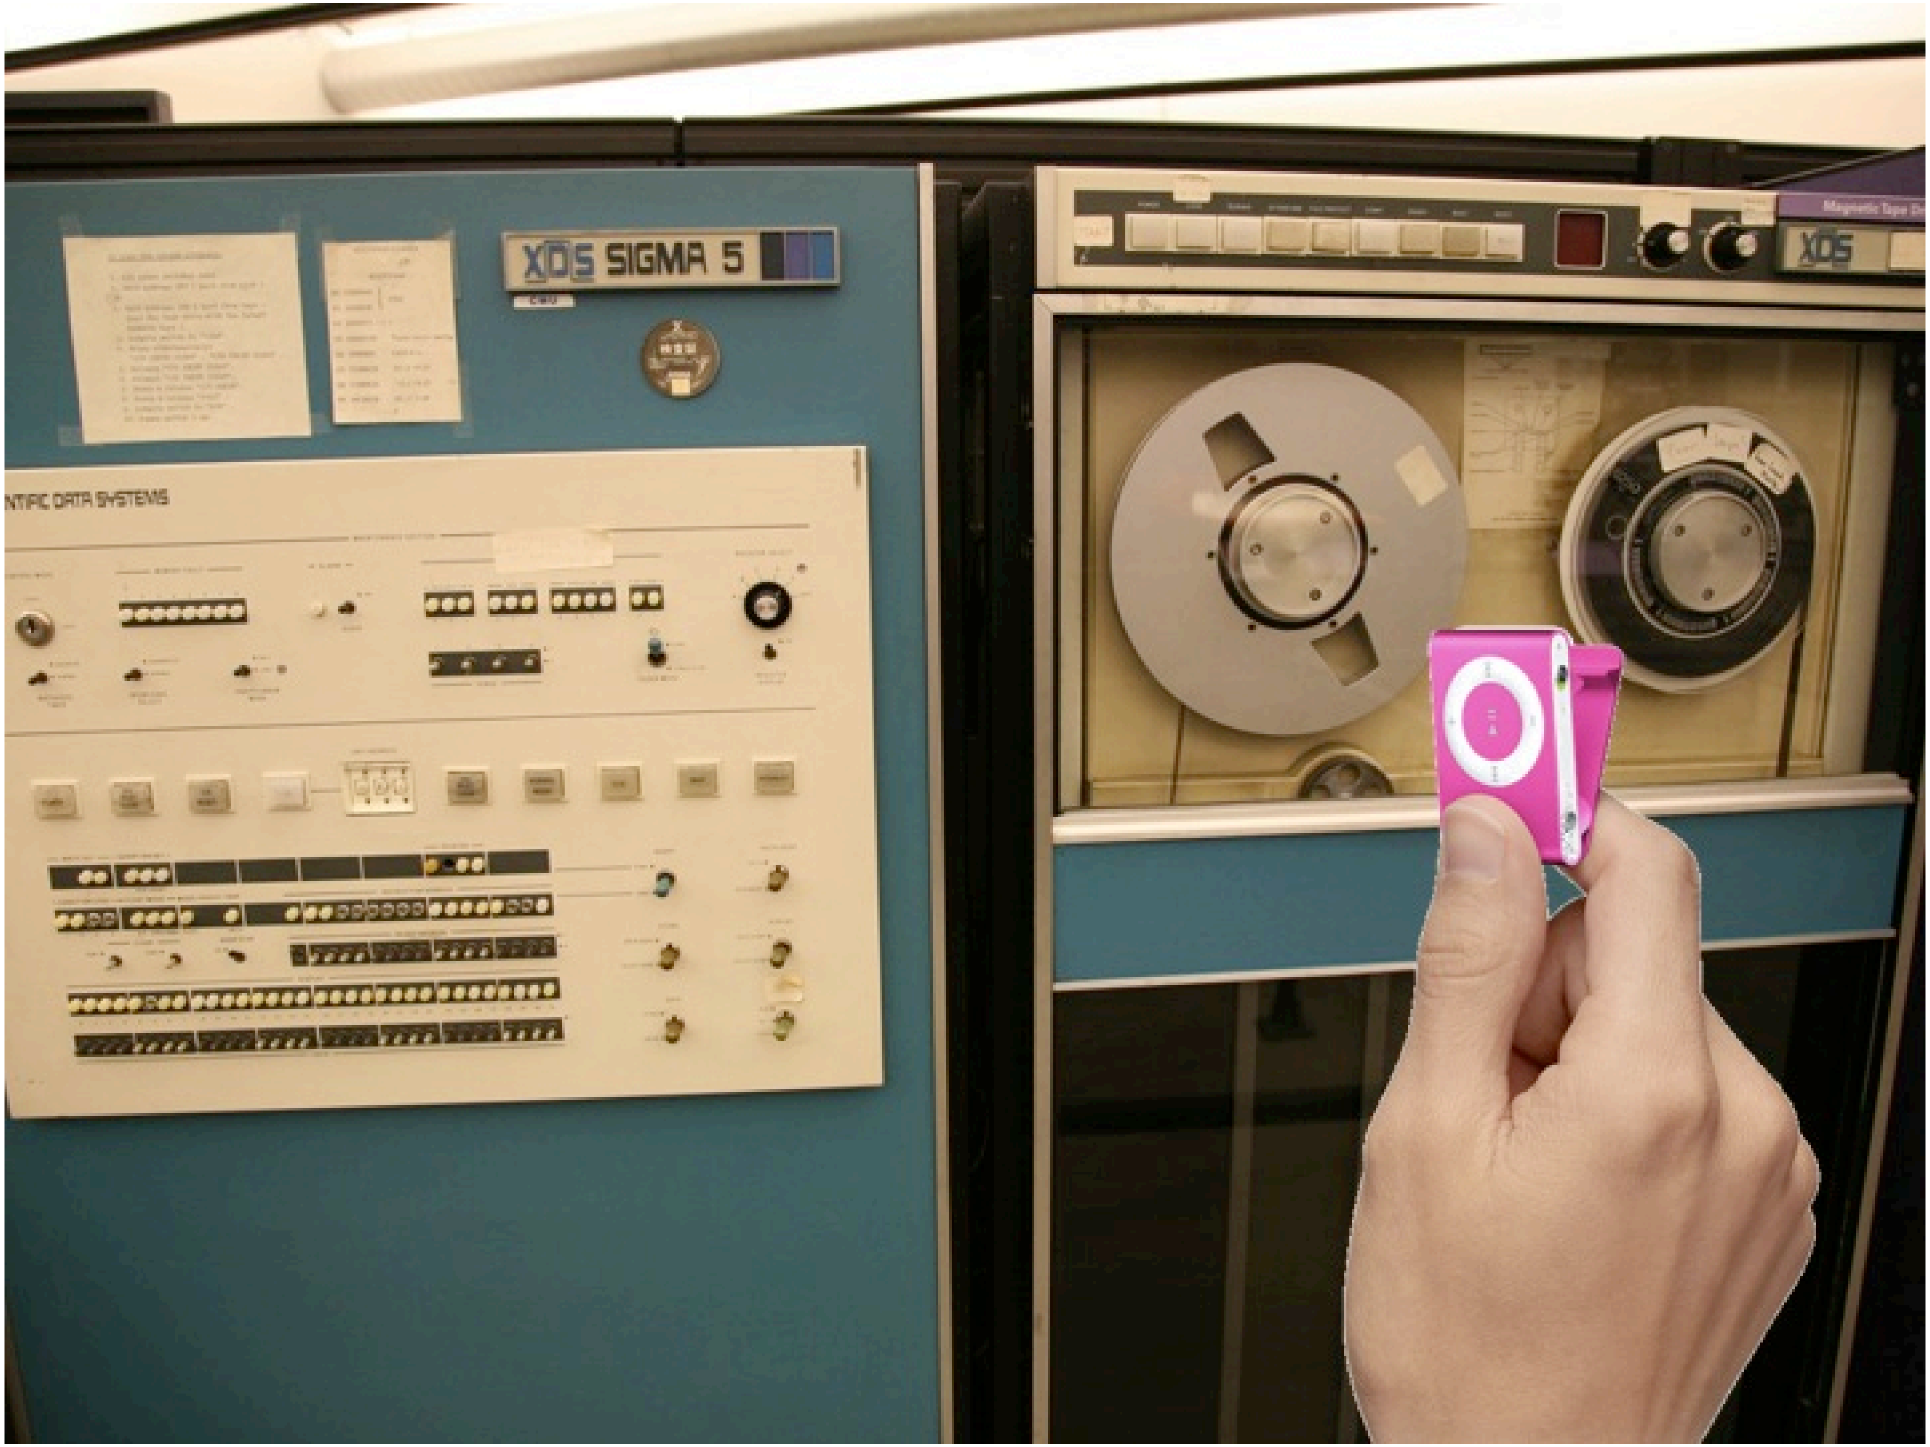
XDS SIGMA 5
NTIFIC DATA SYSTEMS
XDS

ATNF Parkes Swinburne Recorder

- Combination of FPGA and software
  - PDFB3 implements polyphase filterbank
  - 16 processing nodes receive sub-band as UDP stream
- Real-time and/or offline data reduction
  - record data to disk at 1.6 GB/s for 2.5 hours
- Remote control and monitoring
  - web-based interface

SWINBURNE UNIVERSITY OF TECHNOLOGY

Parkes Digital Filterbank

- Designed/developed at ATNF
  - 2 x Compact Array Broad-band (CABB) board
  - up to 2048 channel polyphase filterbank
  - real-time RFI mitigation
- CABB = 2GHz correlator
  - modularity of FPGA design
  - hardware re-use

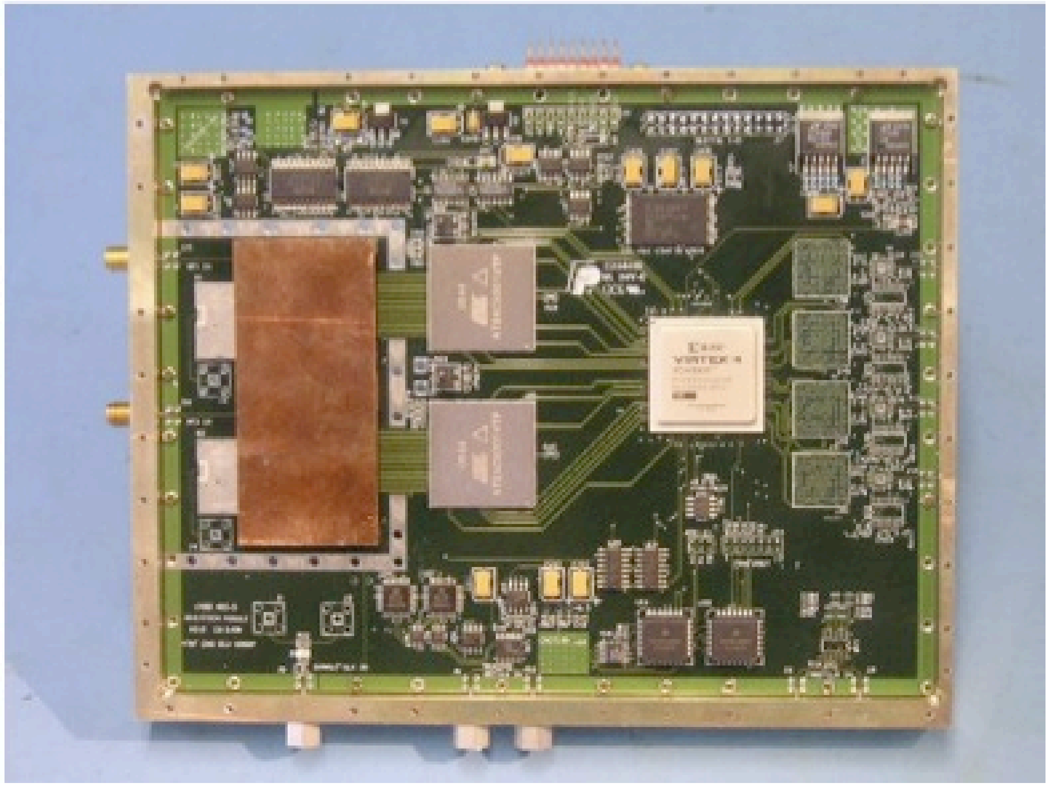

SWINBURNE UNIVERSITY OF TECHNOLOGY

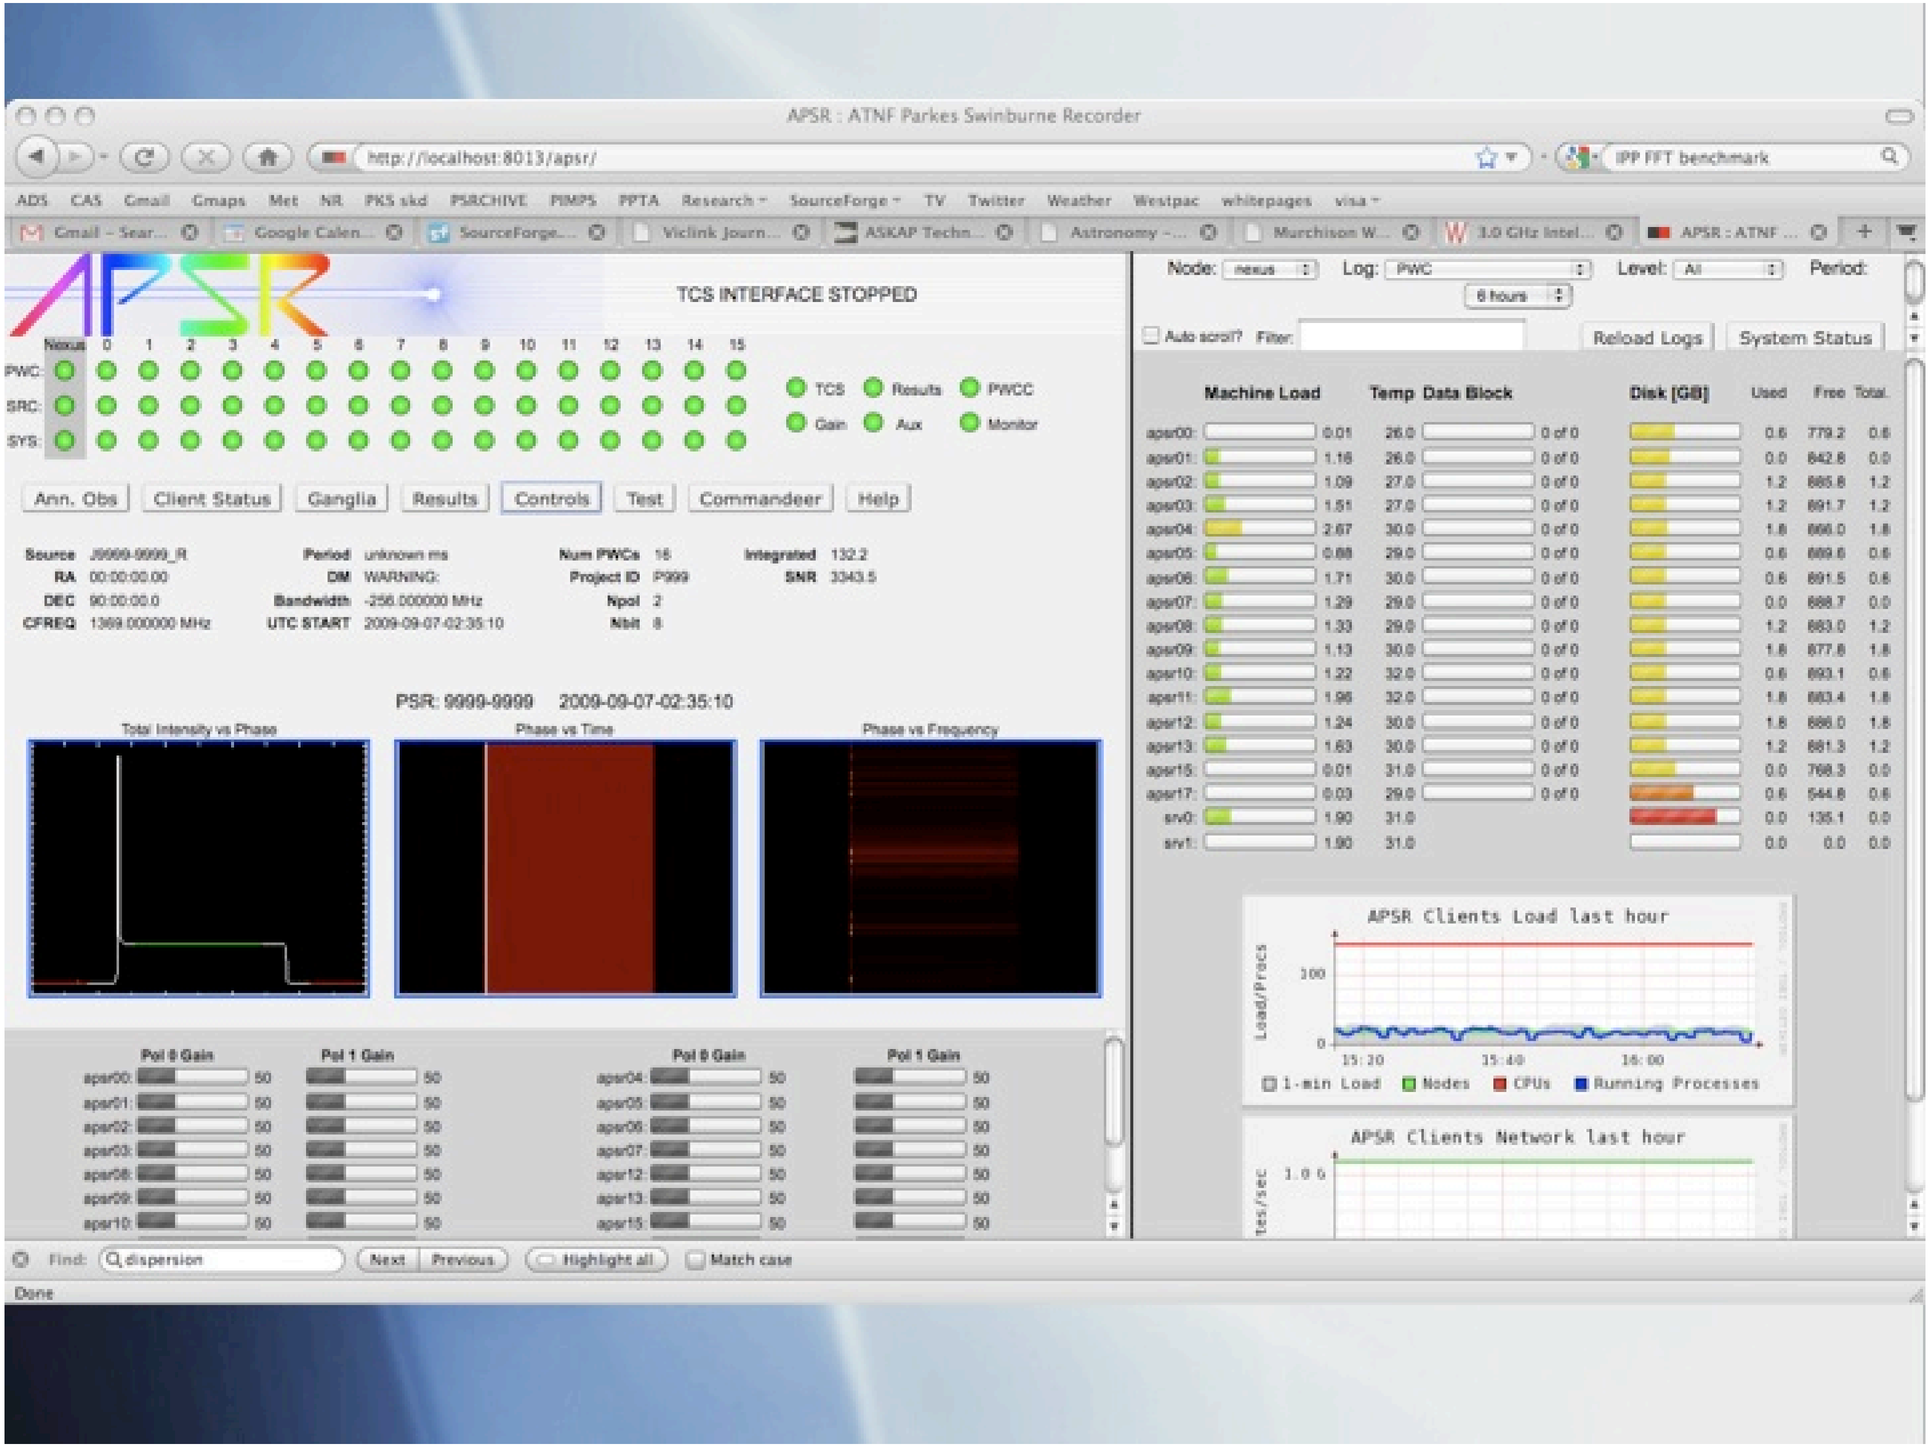

APSR : ATNF Parkes Swinburne Recorder
http://localhost:8013/apsr/
TCS INTERFACE STOPPED
PWC:
SRC:
SYS:
TCS
Results
PWCC
Gain
Aux
Monitor
Ann. Obs
Client Status
Ganglia
Results
Controls
Test
Commandeer
Help
Source J9999-9999_R
RA 00:00:00.00
DEC 90:00:00.0
CFREQ 1369.000000 MHz
Period unknown ms
DM WARNING:
Bandwidth -256.000000 MHz
UTC START 2009-09-07-02:35:10
Num PWCs 16
Project ID P999
Npol 2
Nbit 8
Integrated 132.2
SNR 3343.5
PSR: 9999-9999 2009-09-07-02:35:10
Total Intensity vs Phase
Phase vs Time
Phase vs Frequency
Pol 0 Gain
Pol 1 Gain
Node:
Log:
Level:
Period:
Auto scroll?
Filter:
Reload Logs
System Status
Machine Load
Temp
Data Block
Disk [GB]
Used
Free
Total
APSR Clients Load last hour
Load/Procs
1-min Load
Nodes
CPUs
Running Processes
APSR Clients Network last hour
Find:
Next
Previous
Highlight all
Match case
Done

APSR Features

- Phase-coherent dispersion removal (up to 1024 MHz)
- Impulsive interference excision
  - RFI, lightning, etc.
- Single-pulse capability
  - with real-time calibration and/or giant pulse selection
- Fold multiple pulsars simultaneously
  - globular clusters, binary pulsar

SWIN BUR NE SWINBURNE UNIVERSITY OF TECHNOLOGY

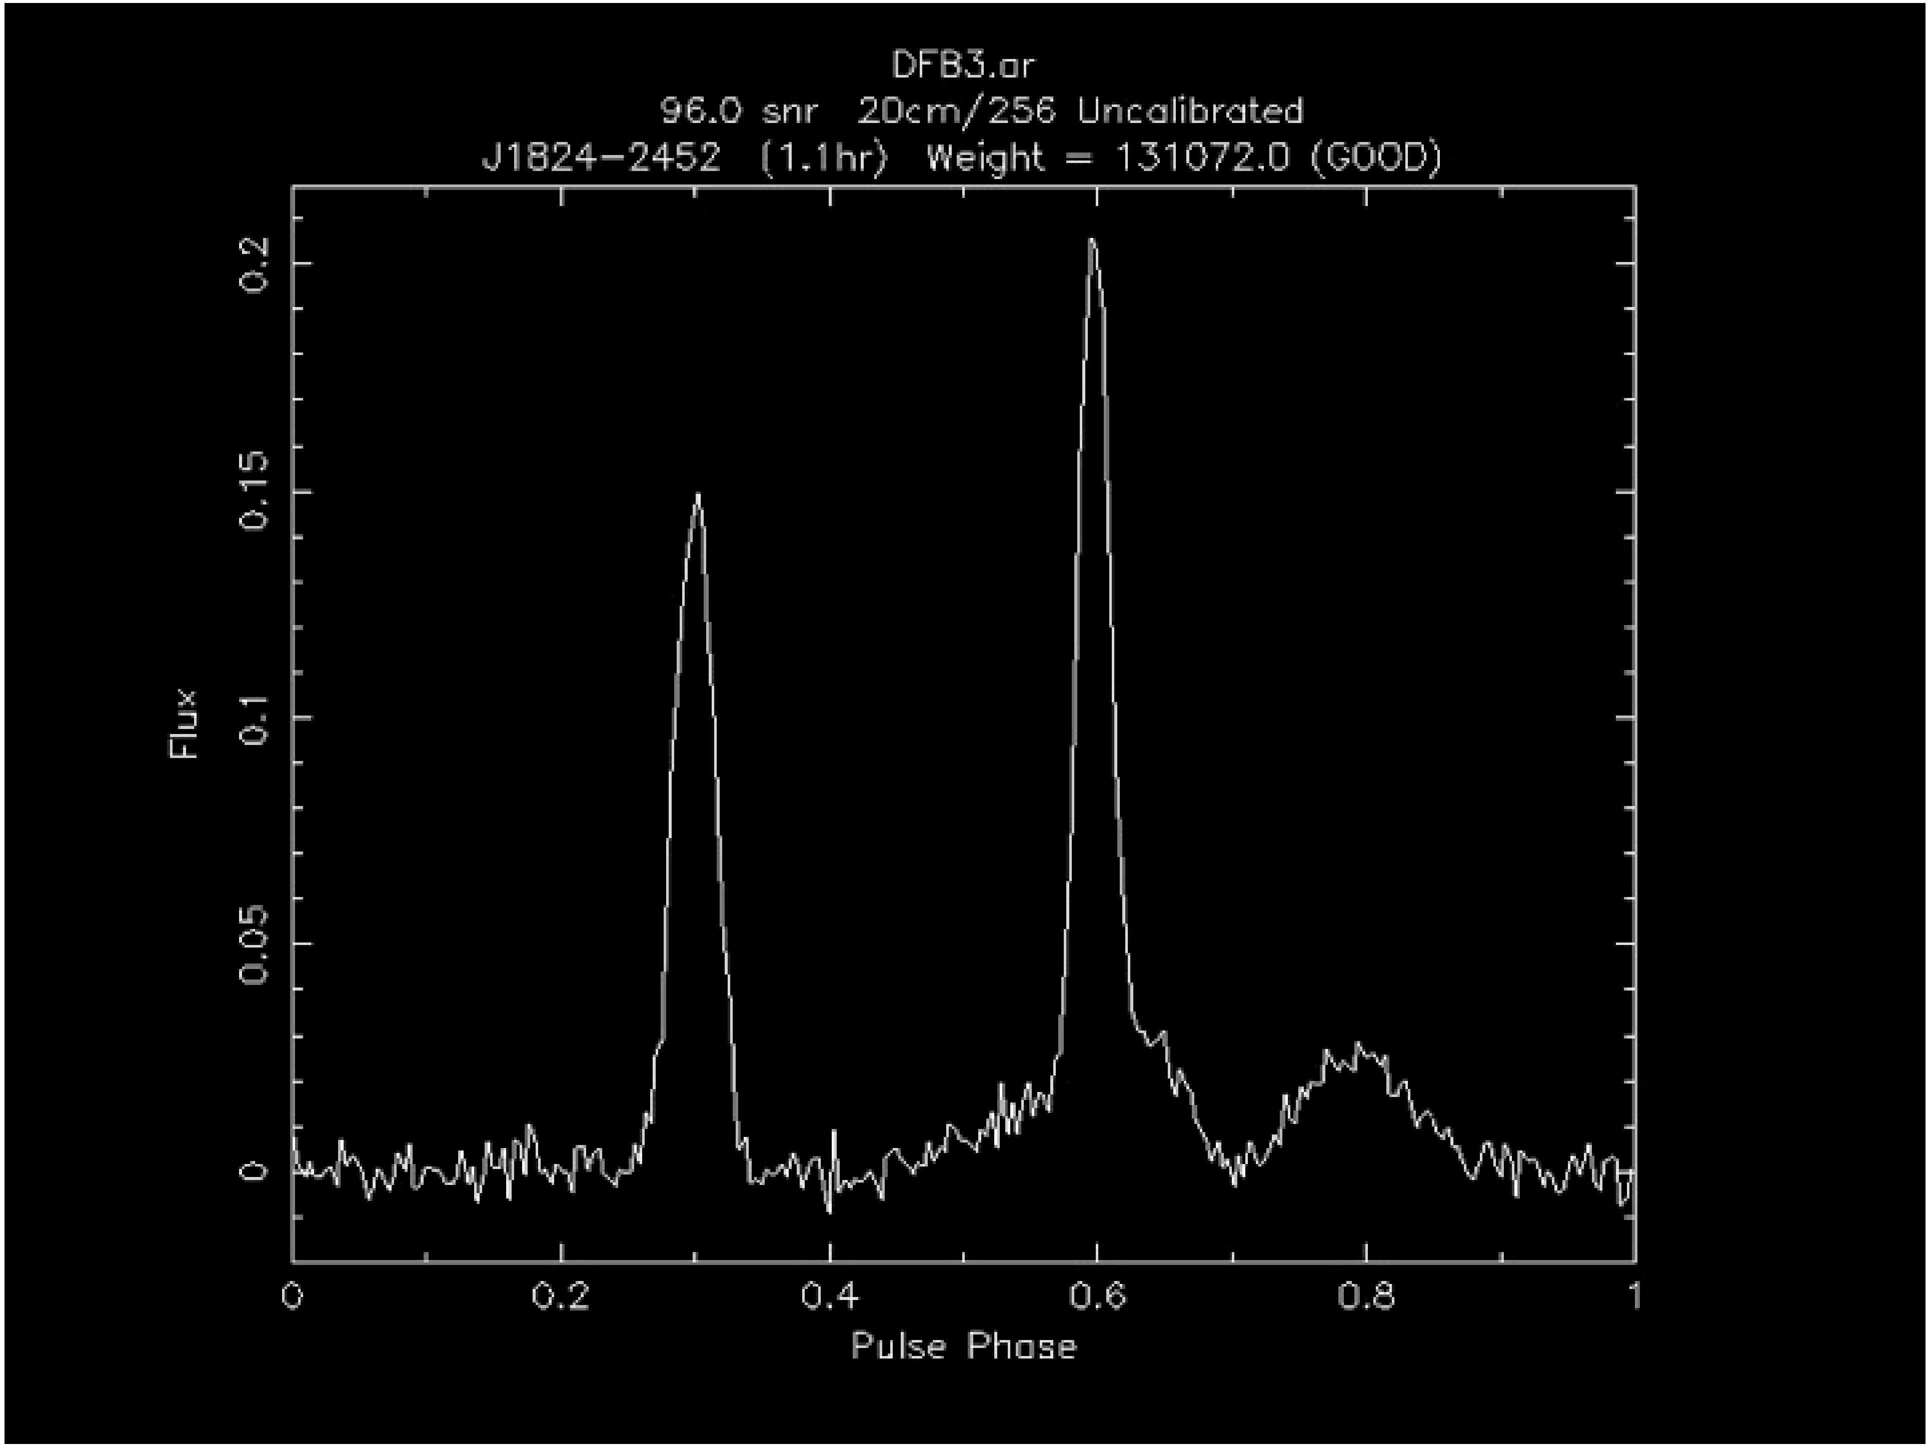

DFB3.ar
96.0 snr 20cm/256 Uncalibrated
J1824-2452 (1.1hr) Weight = 131072.0 (GOOD)
Flux
0
0.05
0.1
0.15
0.2
0
0.2
0.4
0.6
0.8
1
Pulse Phase

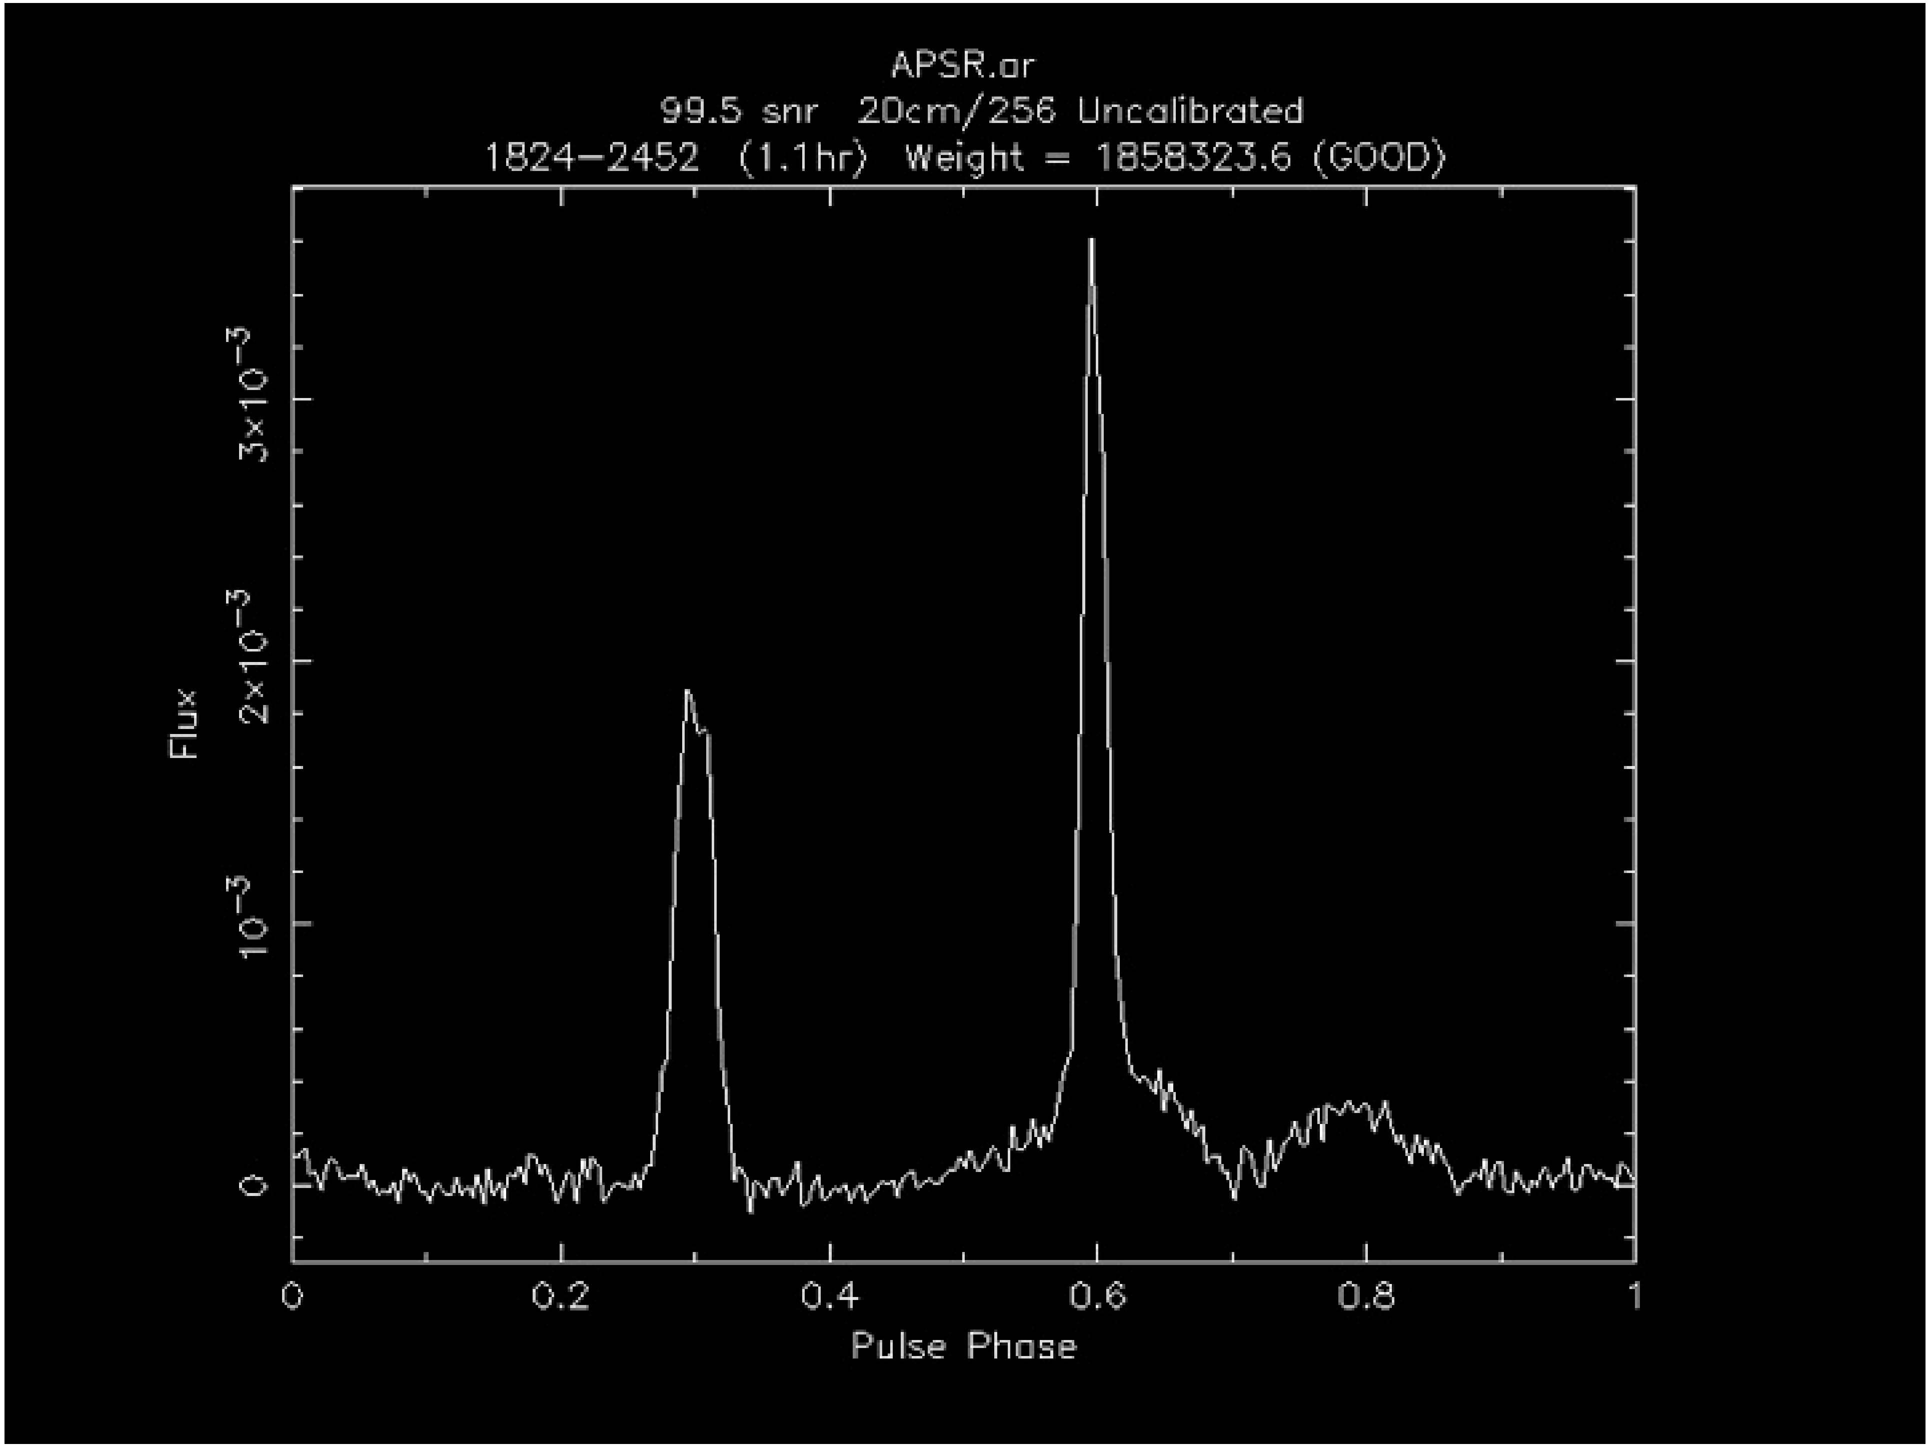
APSR.ar
99.5 snr 20cm/256 Uncalibrated
1824-2452 (1.1hr) Weight = 1858323.6 (GOOD)
Flux
0
10⁻³
2×10⁻³
3×10⁻³
0
0.2
0.4
0.6
0.8
1
Pulse Phase

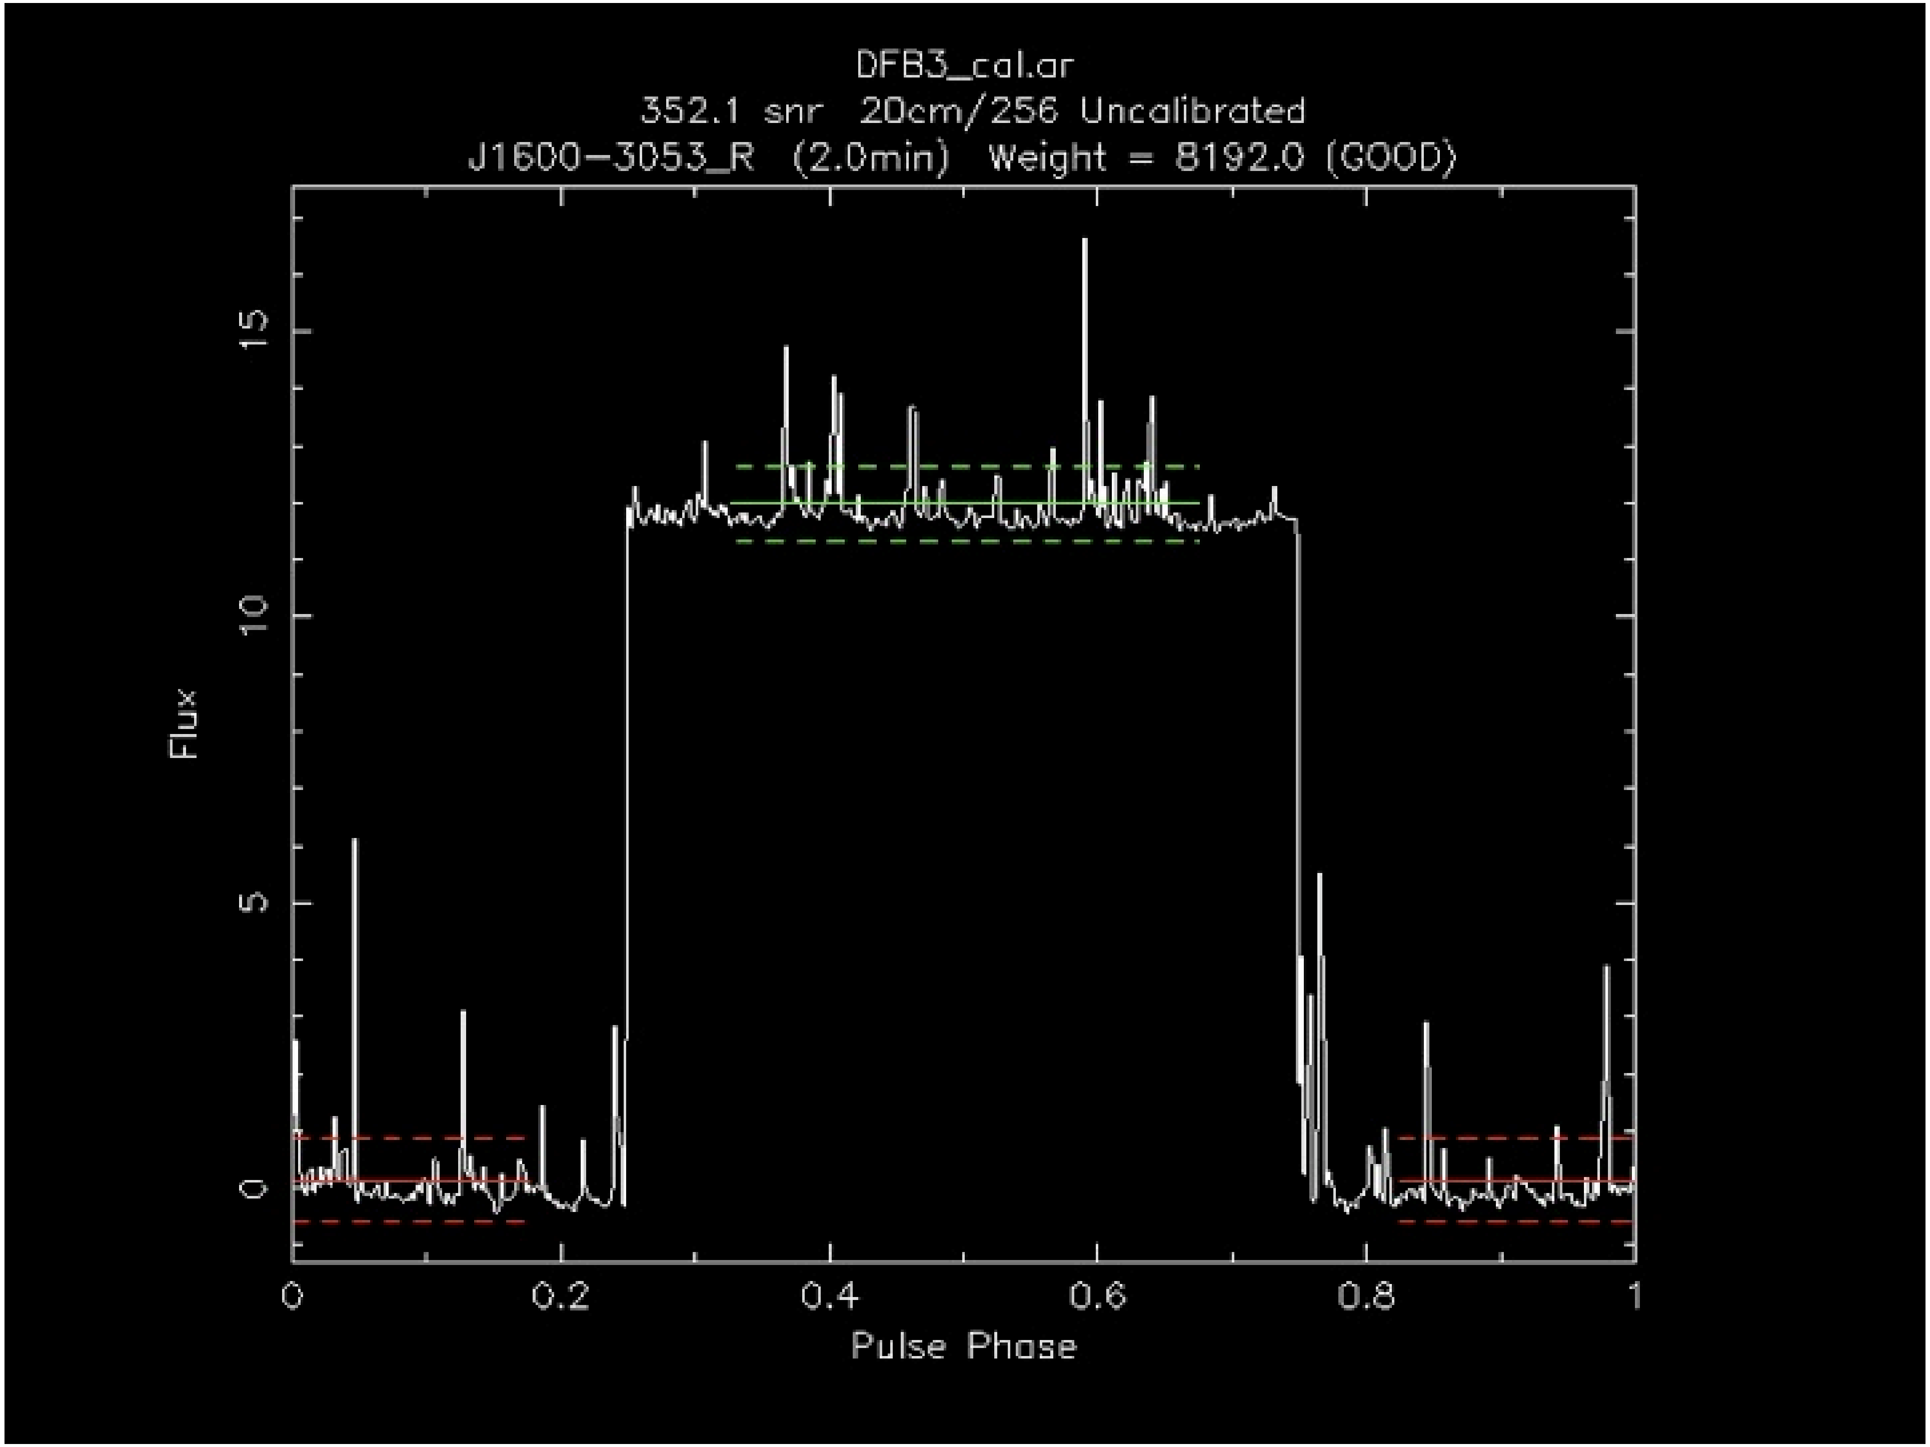

DFB3_cal.ar
352.1 snr 20cm/256 Uncalibrated
J1600-3053_R (2.0min) Weight = 8192.0 (GOOD)
Flux
Pulse Phase
0
5
10
15
0
0.2
0.4
0.6
0.8
1

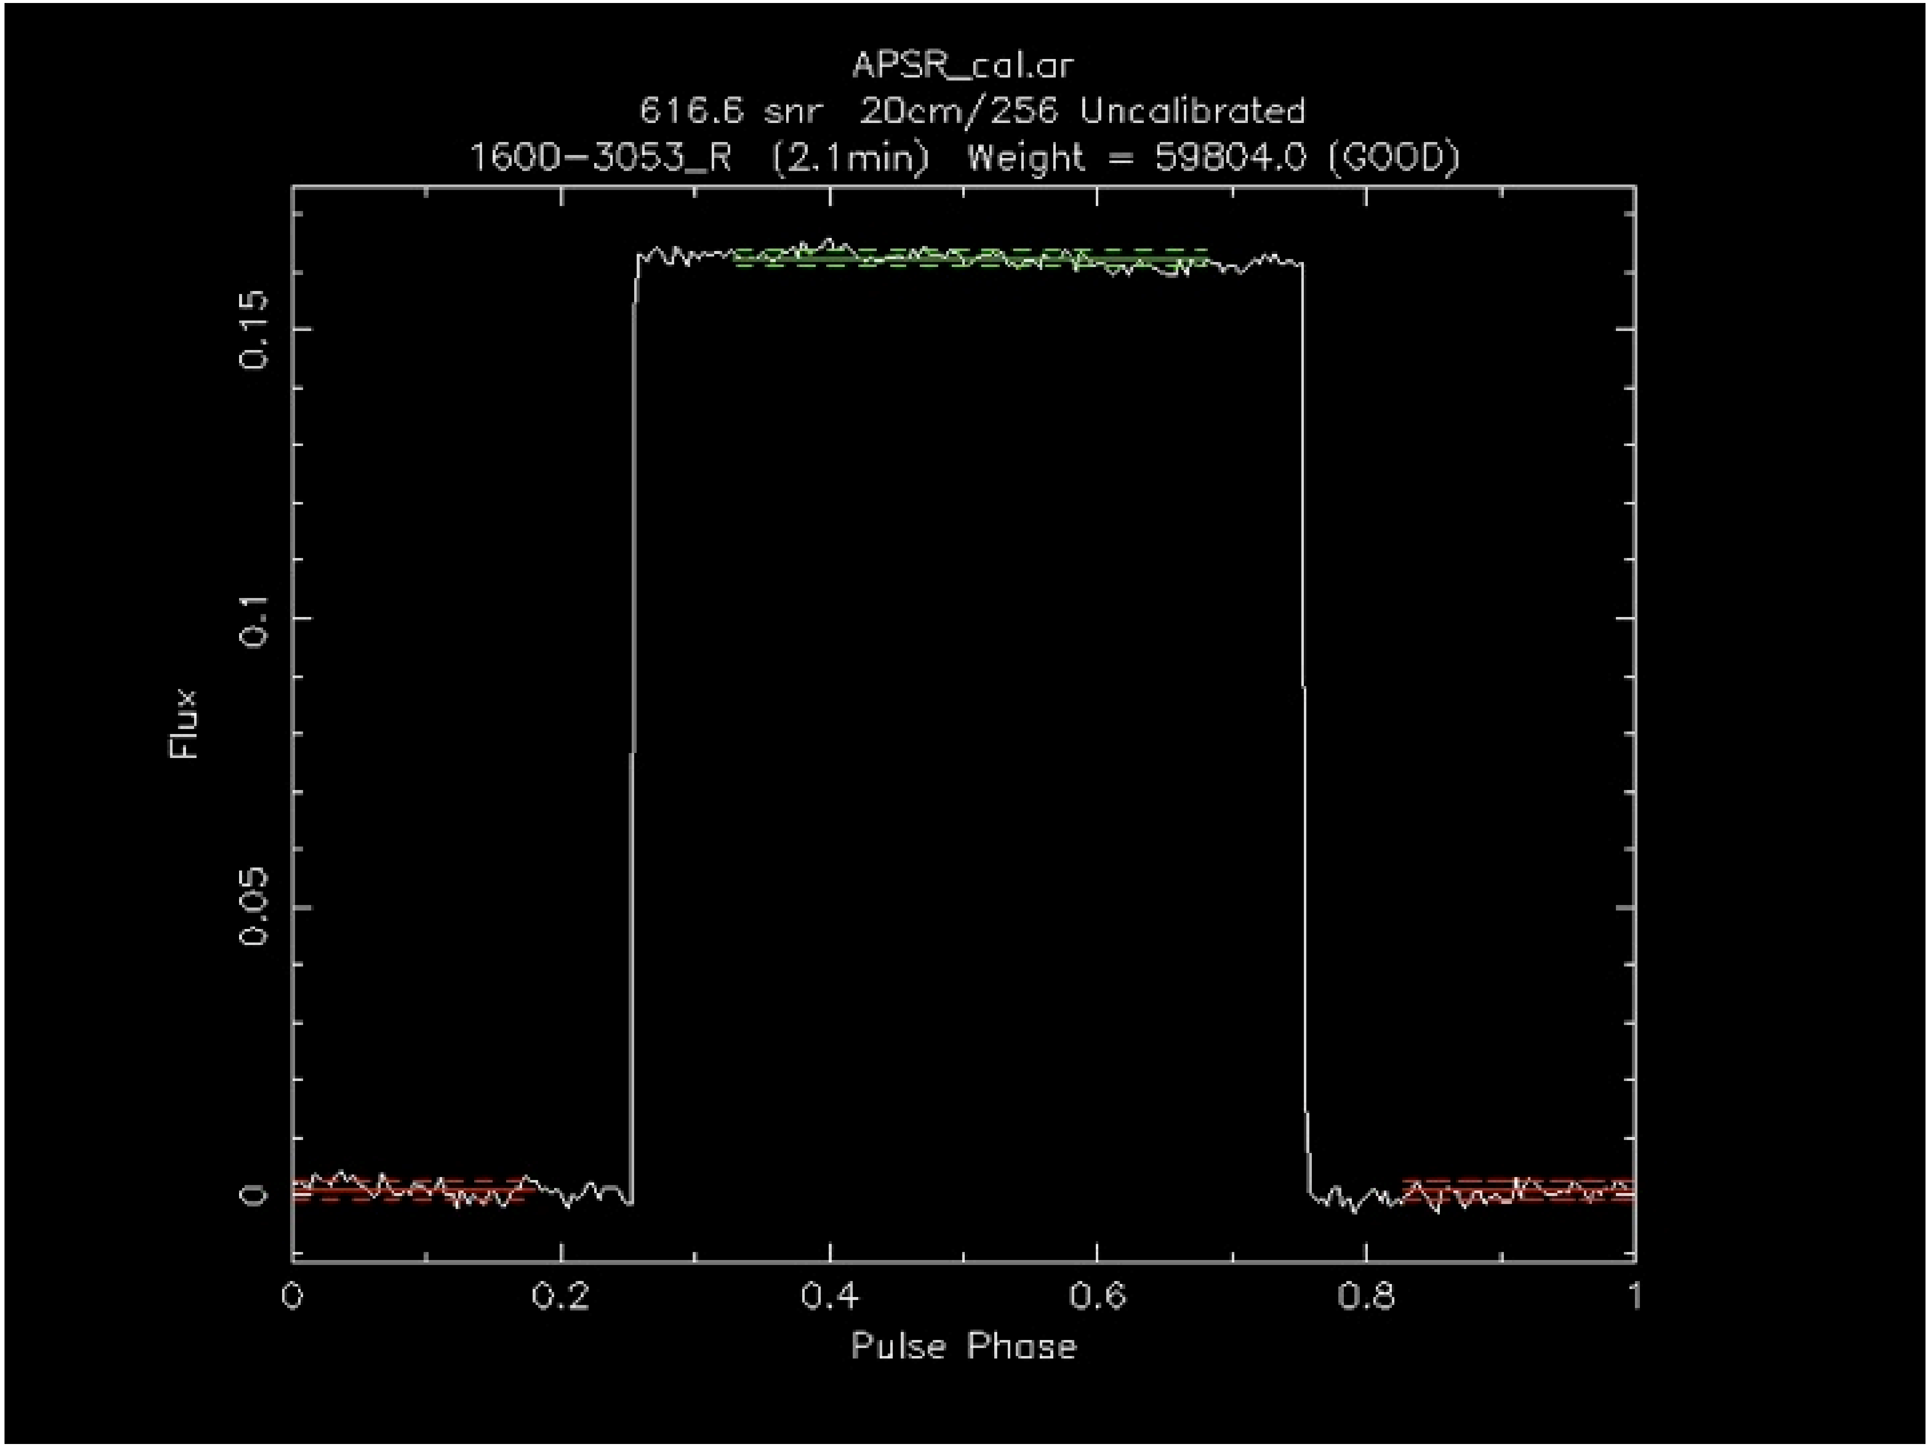

APSR_cal.ar
616.6 snr 20cm/256 Uncalibrated
1600-3053_R (2.1min) Weight = 59804.0 (GOOD)
Flux
Pulse Phase
0
0.05
0.1
0.15
0
0.2
0.4
0.6
0.8
1

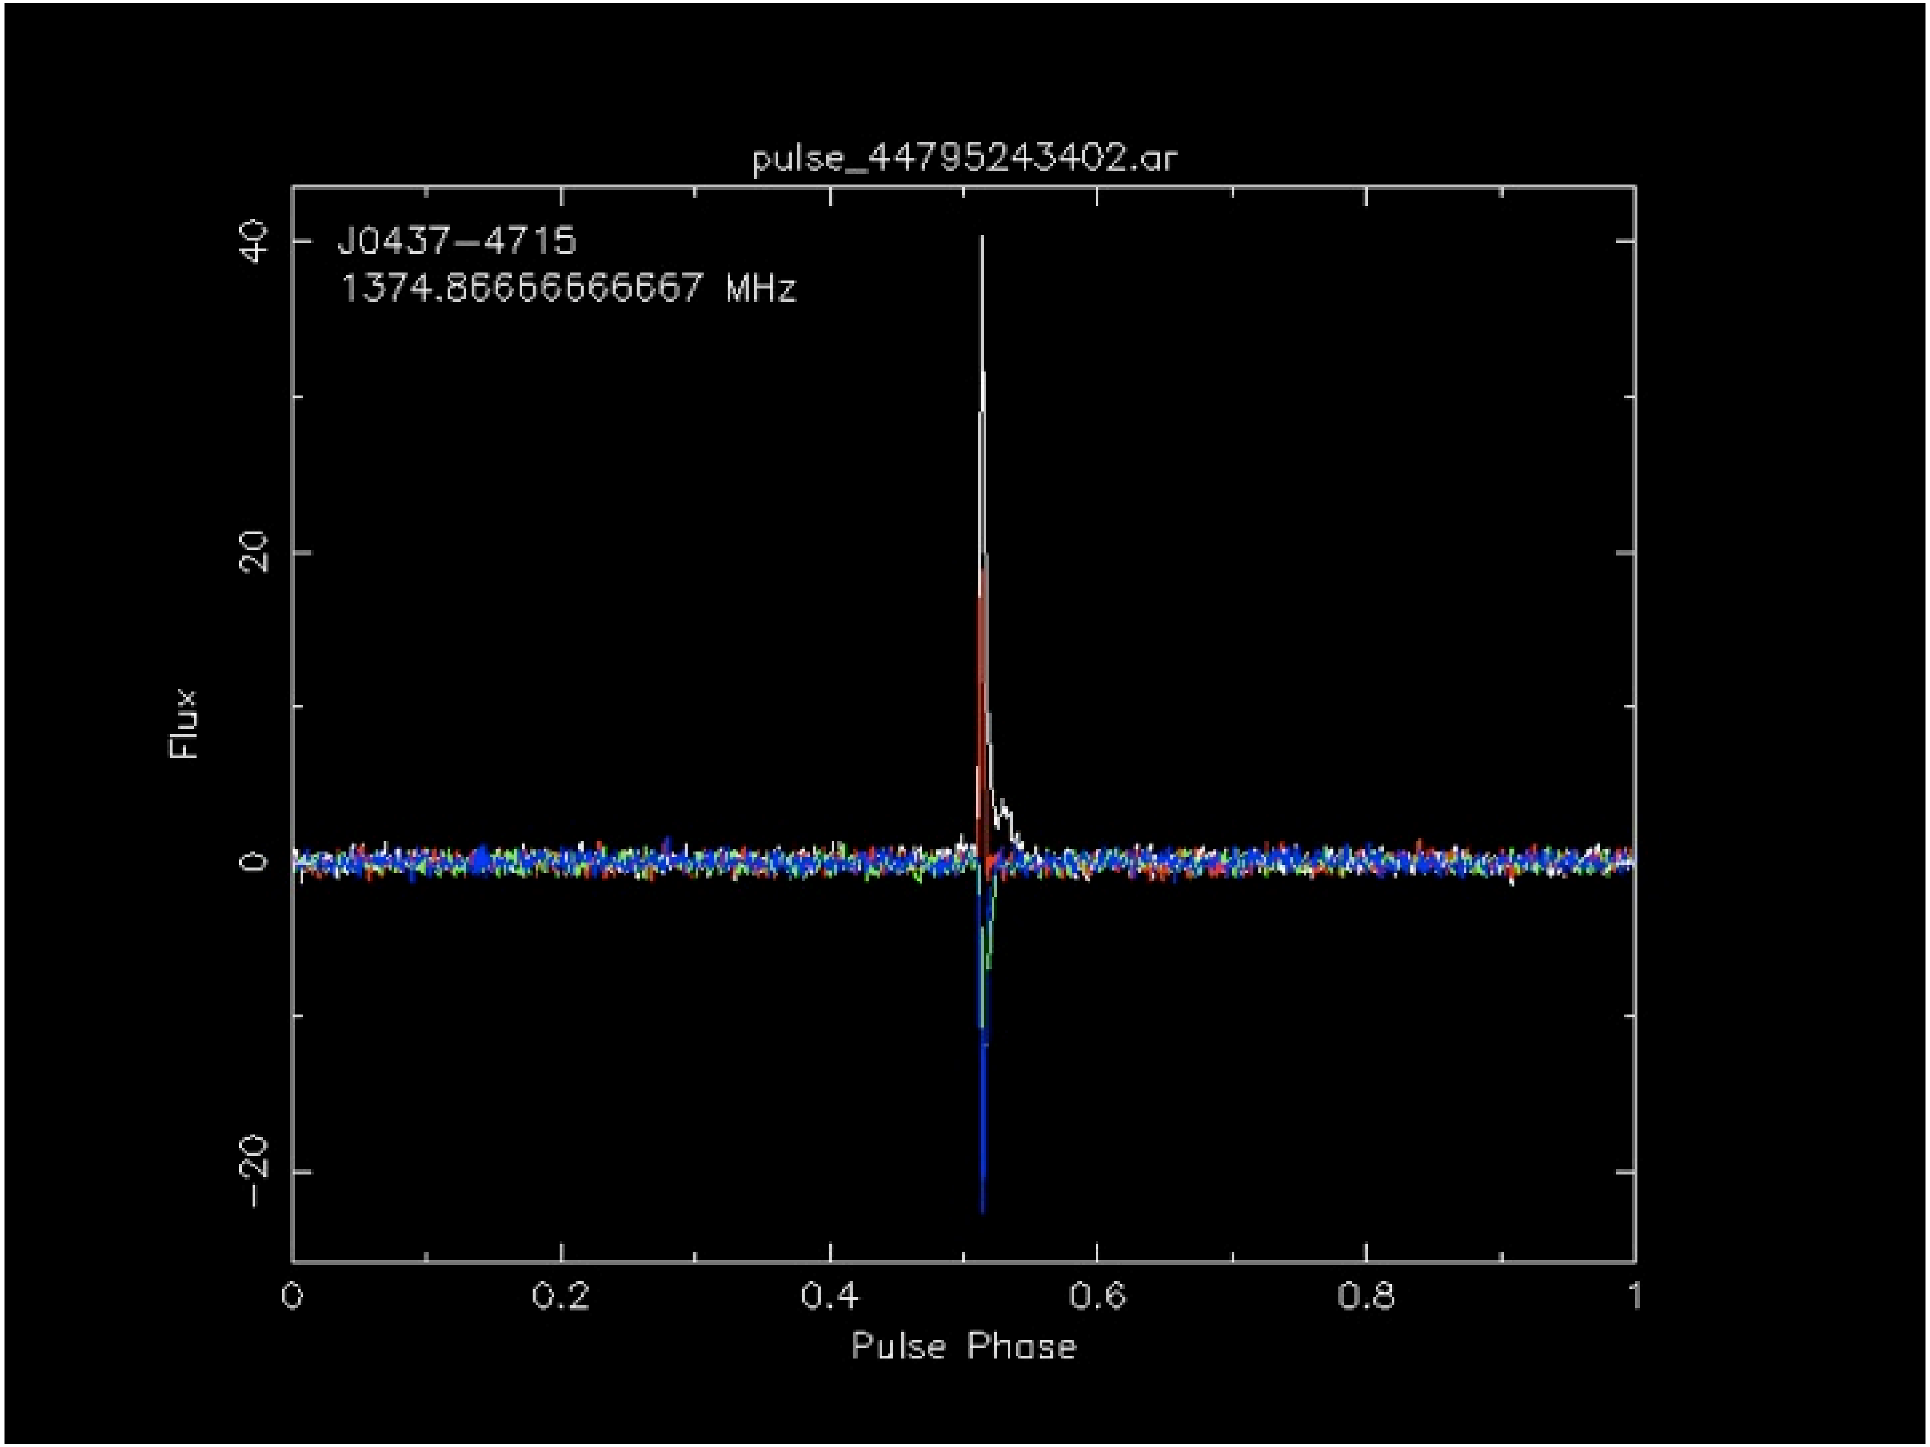
pulse_44795243402.ar
J0437-4715
1374.86666666667 MHz
Flux
Pulse Phase
40
20
0
-20
0
0.2
0.4
0.6
0.8
1

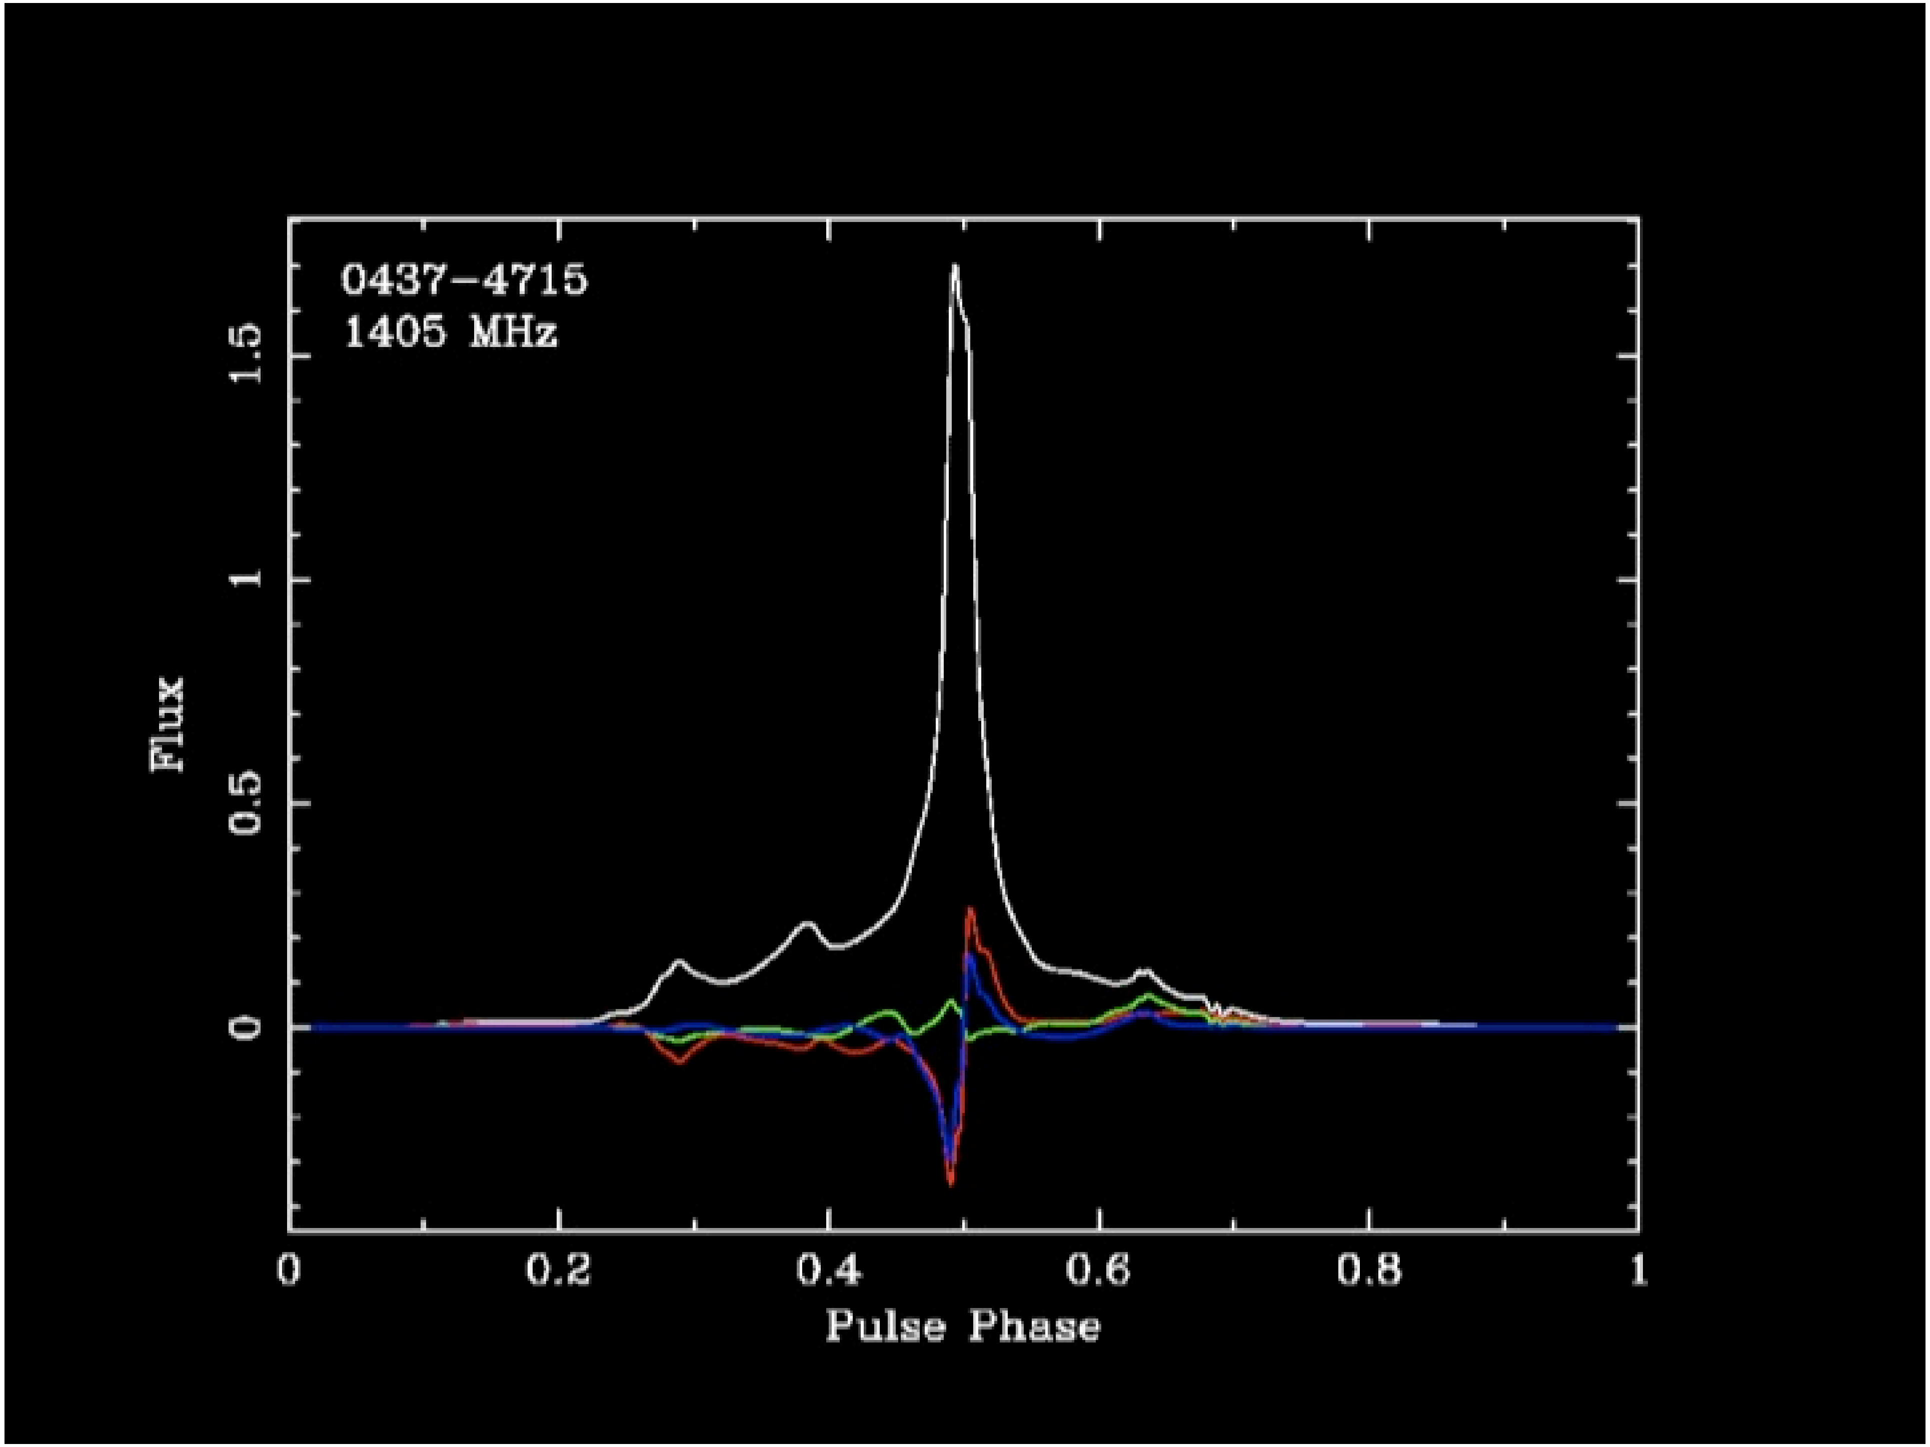

0437-4715
1405 MHz
Flux
Pulse Phase
0
0.5
1
1.5
0
0.2
0.4
0.6
0.8
1

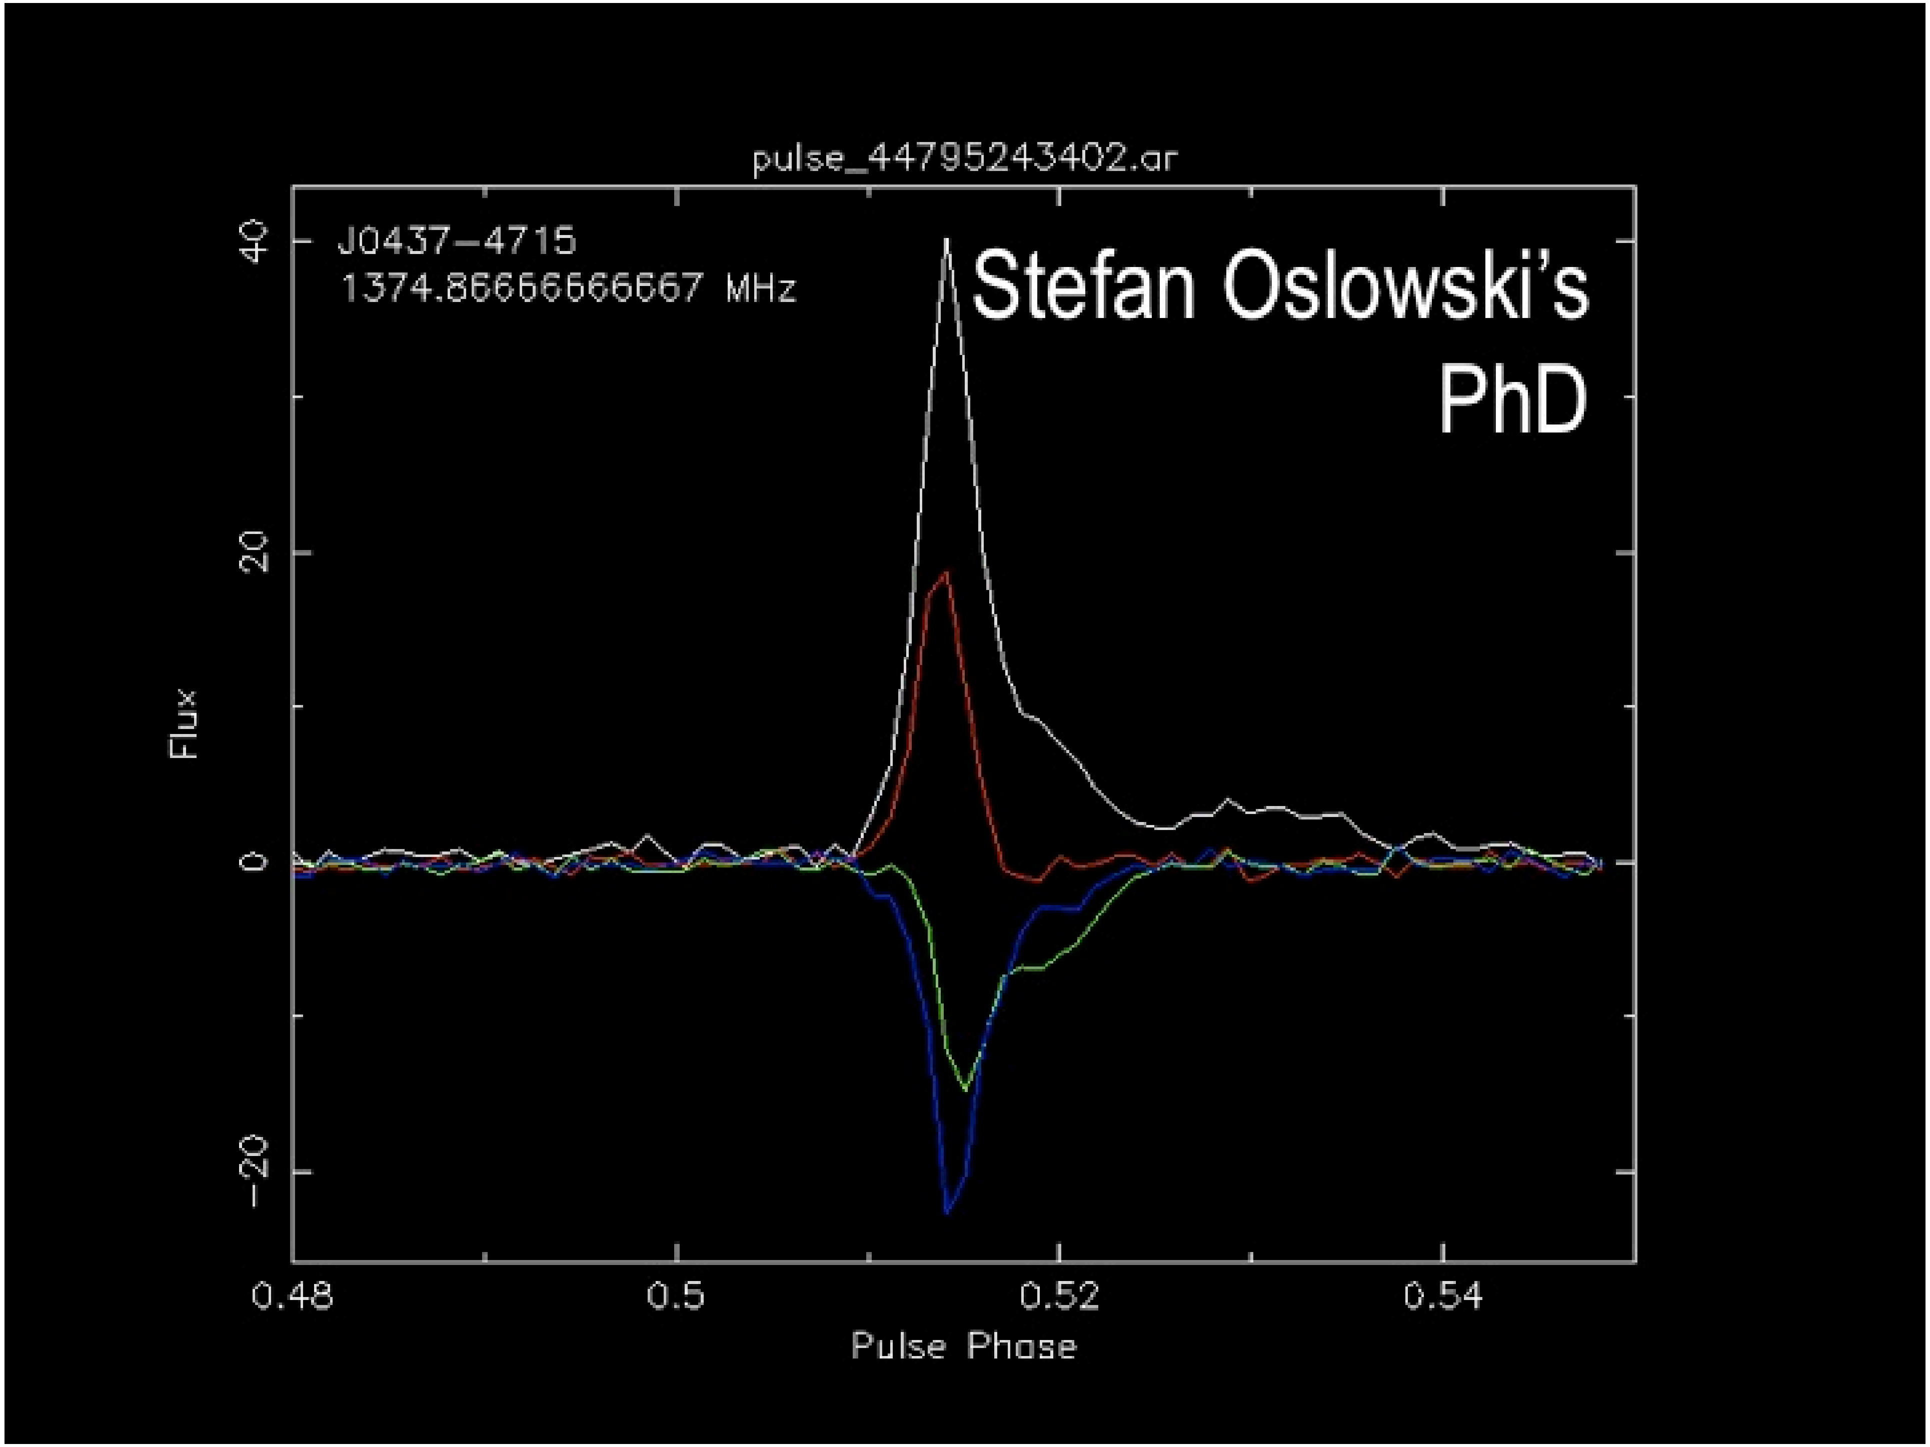

pulse_44795243402.ar
J0437-4715
1374.86666666667 MHz
Stefan Oslowski's
PhD
40
20
0
-20
Flux
0.48
0.5
0.52
0.54
Pulse Phase

Pulsar Surveys

- Computationally intensive search over:
  - period, P
  - dispersion measure, DM
  - acceleration, a
- Point at patch of sky and record:
  - sensitivity is a function of period and DM

Parkes Multibeam Surveys

- Discovered 850+ pulsars
  - Double pulsar
  - NS+NS and NS+WD binaries
  - Energetic and magnetar-like pulsars
  - Millisecond pulsars (for high-precision timing)

- 13 beams x 288 MHz (96 channels x 3 MHz)
- 250 microsecond, 1 bit/sample (624 kB/s)

SWINBURNE UNIVERSITY OF TECHNOLOGY

Berkeley Parkes Swinburne Recorder

- Combination of FPGA and software:
  - ☐ 1024 channels, 64 us, 8bits x 2 poln from iBOB
  - ☐ decimated to 2 bits x total intensity (52 MB/s)
  - ☐ written to DLT S4 tapes (x2)
  - ☐ sent over 1Gb link to Swinburne
  - ☐ processed on supercomputer

- High Time Resolution Universe Survey

SWIN BUR NE SWINBURNE UNIVERSITY OF TECHNOLOGY

Lina Levin's PhD

iBOB

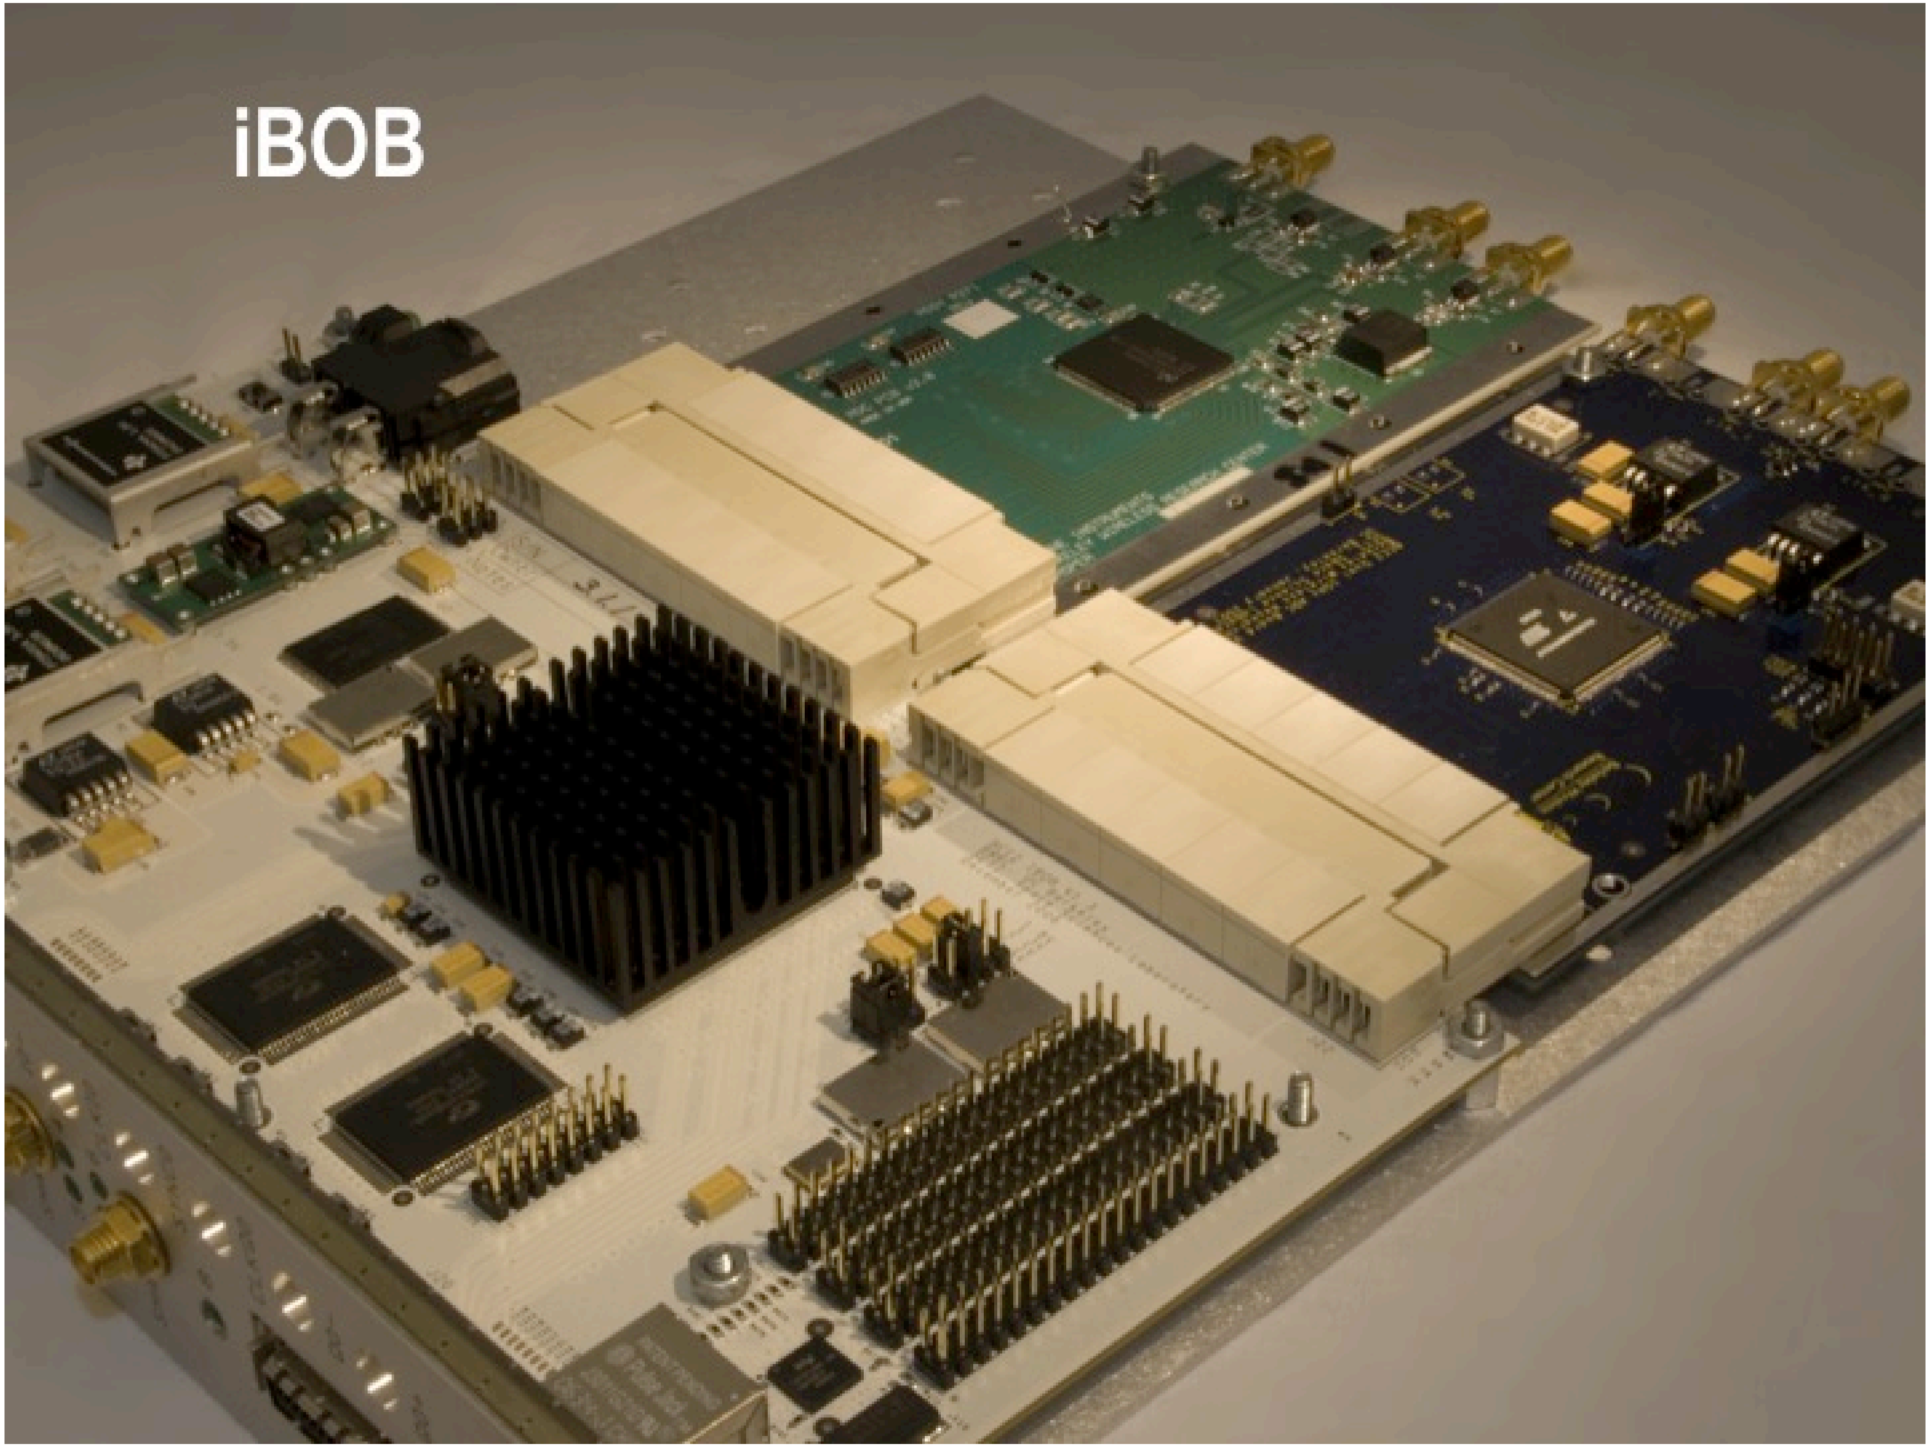

iBOB

- Center for Astronomy Signal Processing and Electronics Research (CASPER) at Berkeley
- FPGA “gateware” libraries
- modular, hardware building blocks
- Latest design:
  - Reconfigurable Open Architecture Computing Hardware (ROACH)

Future: Transient Searches

- search for strong bursts as a function of DM
  - no search over period and acceleration
- real-time detection required to:
  - save raw data for more detailed analysis
  - trigger follow up at other wavelengths
- must differentiate b/w astronomical & terrestrial signals
- need wide field of view

Transient Sources

- Explosive Events
  - ☐ Gamma Ray Bursts, radio supernovae
- Stellar and Planetary Emission
  - ☐ Sun, Jupiter
- Compact Objects
  - ☐ Active Galactic Nuclei (AGN), micro-quasars
- Serendipity
  - ☐ LIGO coincidence, SETI

SWINBURNE UNIVERSITY OF TECHNOLOGY

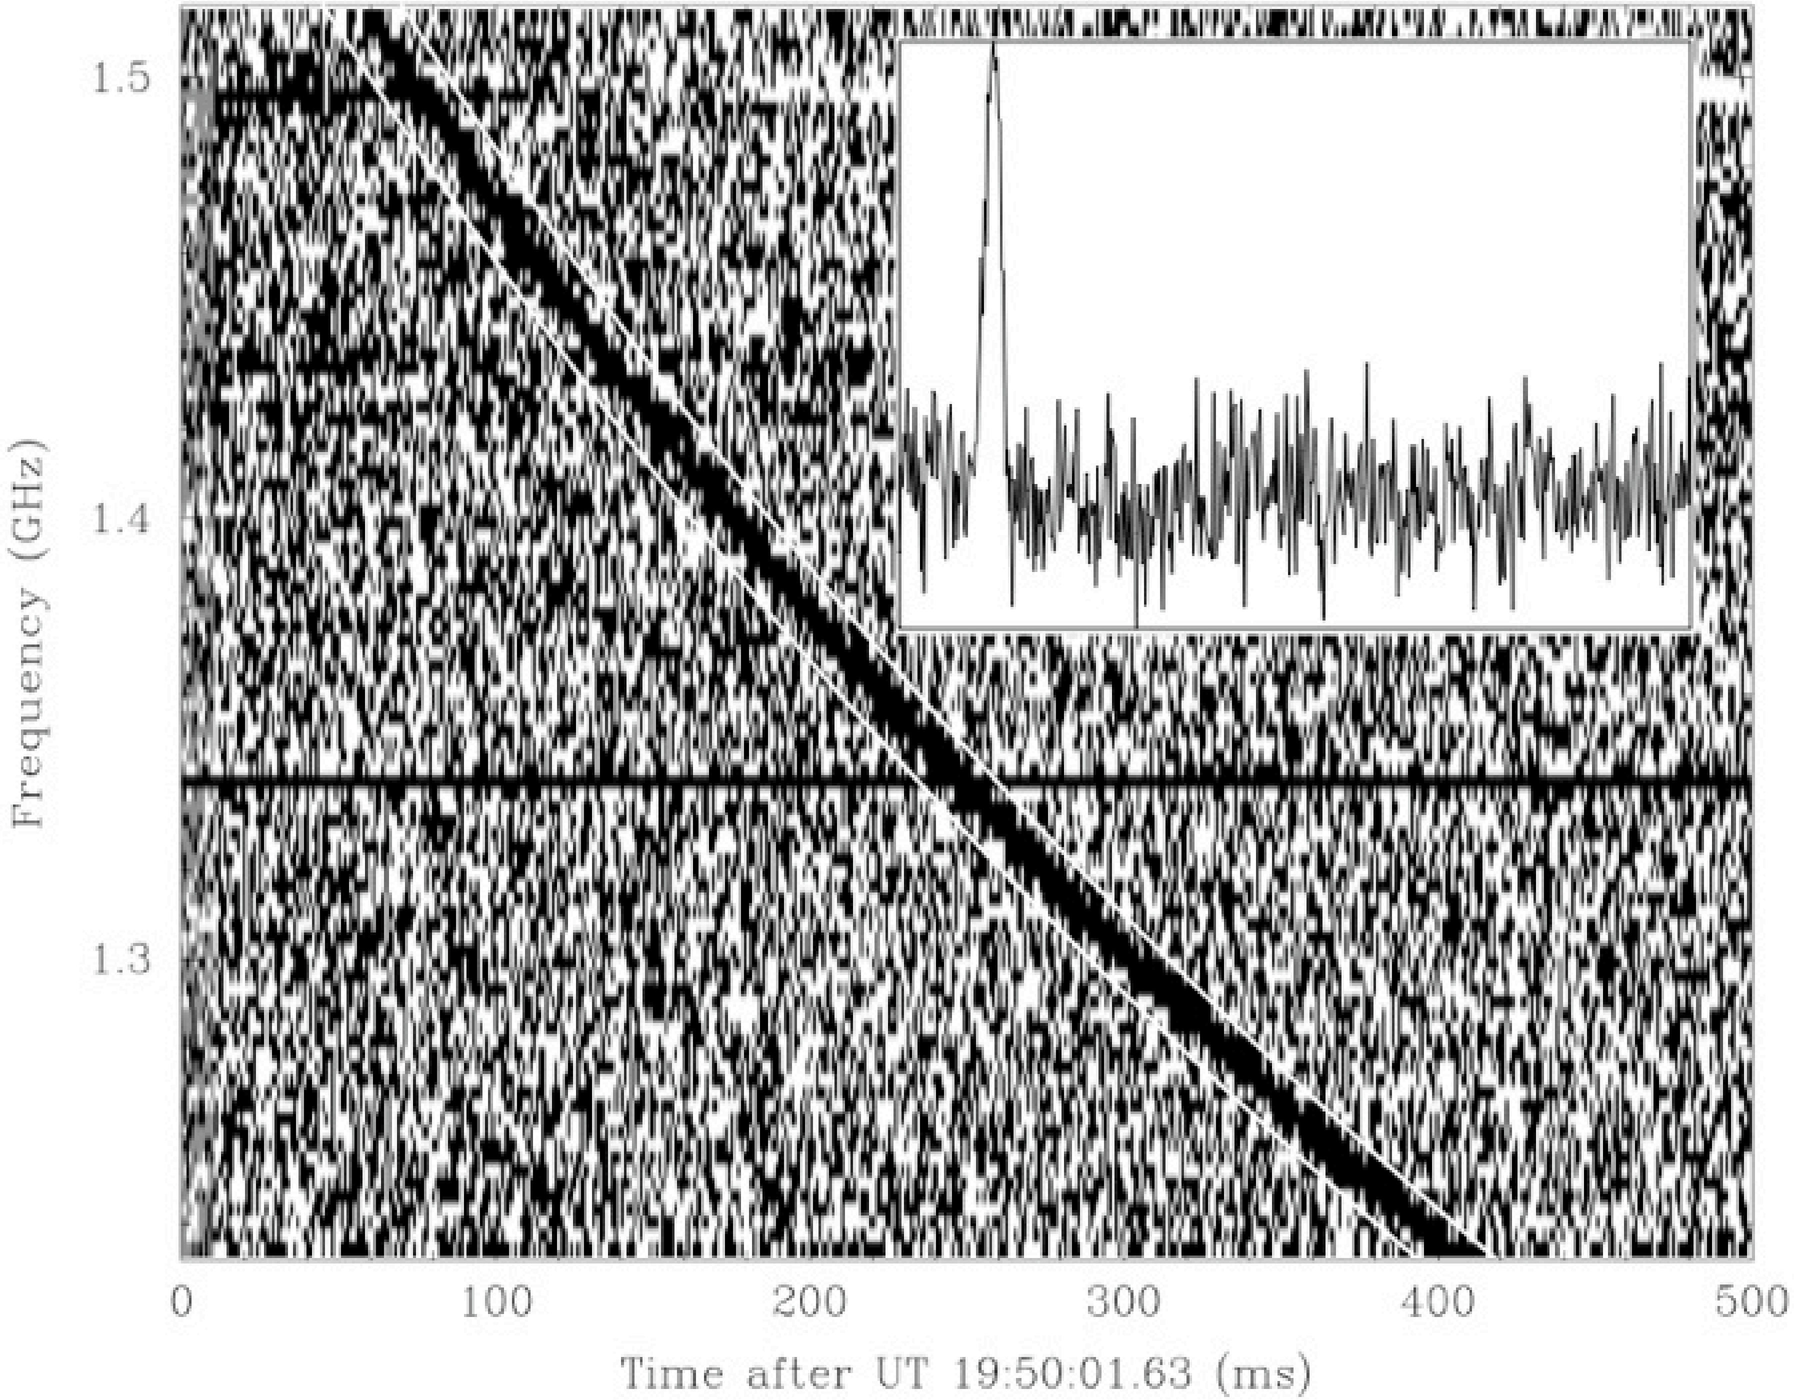

1.5
1.4
1.3
Frequency (GHz)
0
100
200
300
400
500
Time after UT 19:50:01.63 (ms)

Wide Field of View

- Large Number Small Diameter (LNSD)
  - many small antennae
  - large field of view (all sky monitor)

- Focal Plane Arrays
  - fewer, mid-sized antennae
  - phased array at focus produces multiple beams

SWINBURNE UNIVERSITY OF TECHNOLOGY

Tessellate the Primary Beam

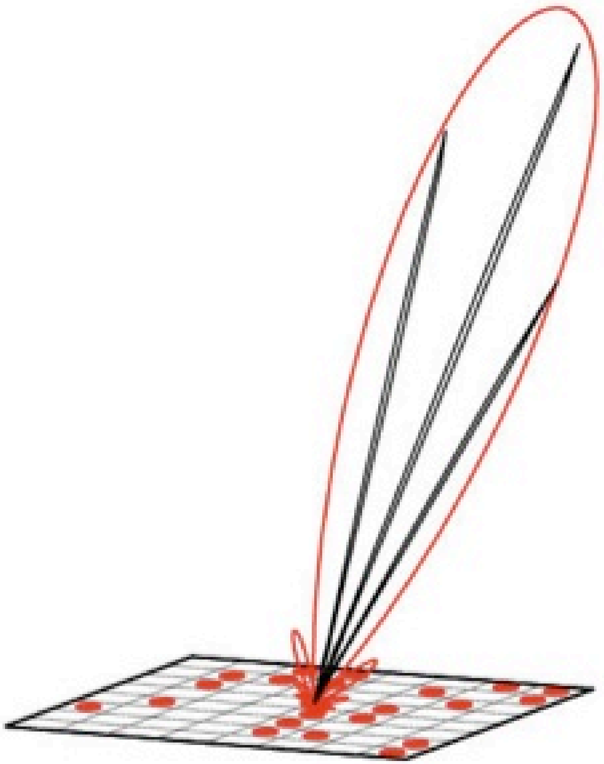

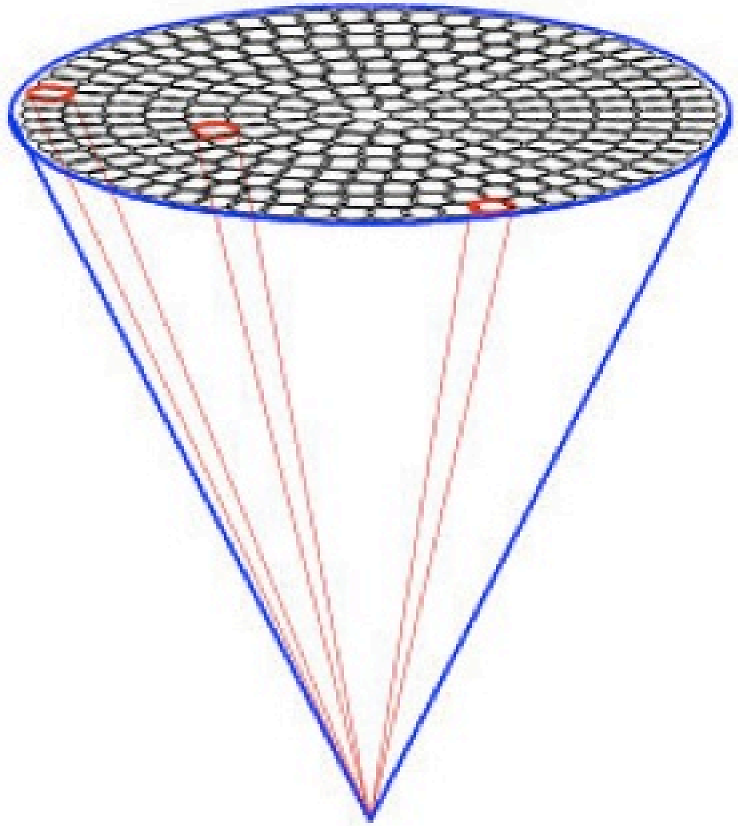

SWINBURNE UNIVERSITY OF TECHNOLOGY

Murchison Widefield Array

- 15-50 degree field of view
- 2.3 - 8.6 arcsec tied-array beam
- 120k - 160k beams!
- How many tied array beams are possible?

SWINBURNE UNIVERSITY OF TECHNOLOGY

Mega beamformer?

- Need: Frequency resolution (B=32MHz N_{FFT}=4096)
  - $O(N_{FFT})$ B/N_{FFT} $N_{tel}N_{pol}$ = ~2 Tflops
- Need: IO
  - 512 Mb/s x N_{tel} = 256 Gb/s (Infiniband 40Gb exists)
- Need: Complex Multiply-ACcumulate
  - B x $N_{tel}N_{pol}$ x 4 = 128 GigaCMACs
- Need: Memory Bandwidth
  - ~ N_{byte}[8] x CMAC ~ 1 TB/sec

per beam!!

SWINBURNE UNIVERSITY OF TECHNOLOGY

Memory Bandwidth

Compute Device	GB/s
Intel Clovertown	~10
Intel Nehalem	~30
nVidia GPU	~100
Intel Larrabee	????

NEED: 1TB/s PER BEAM

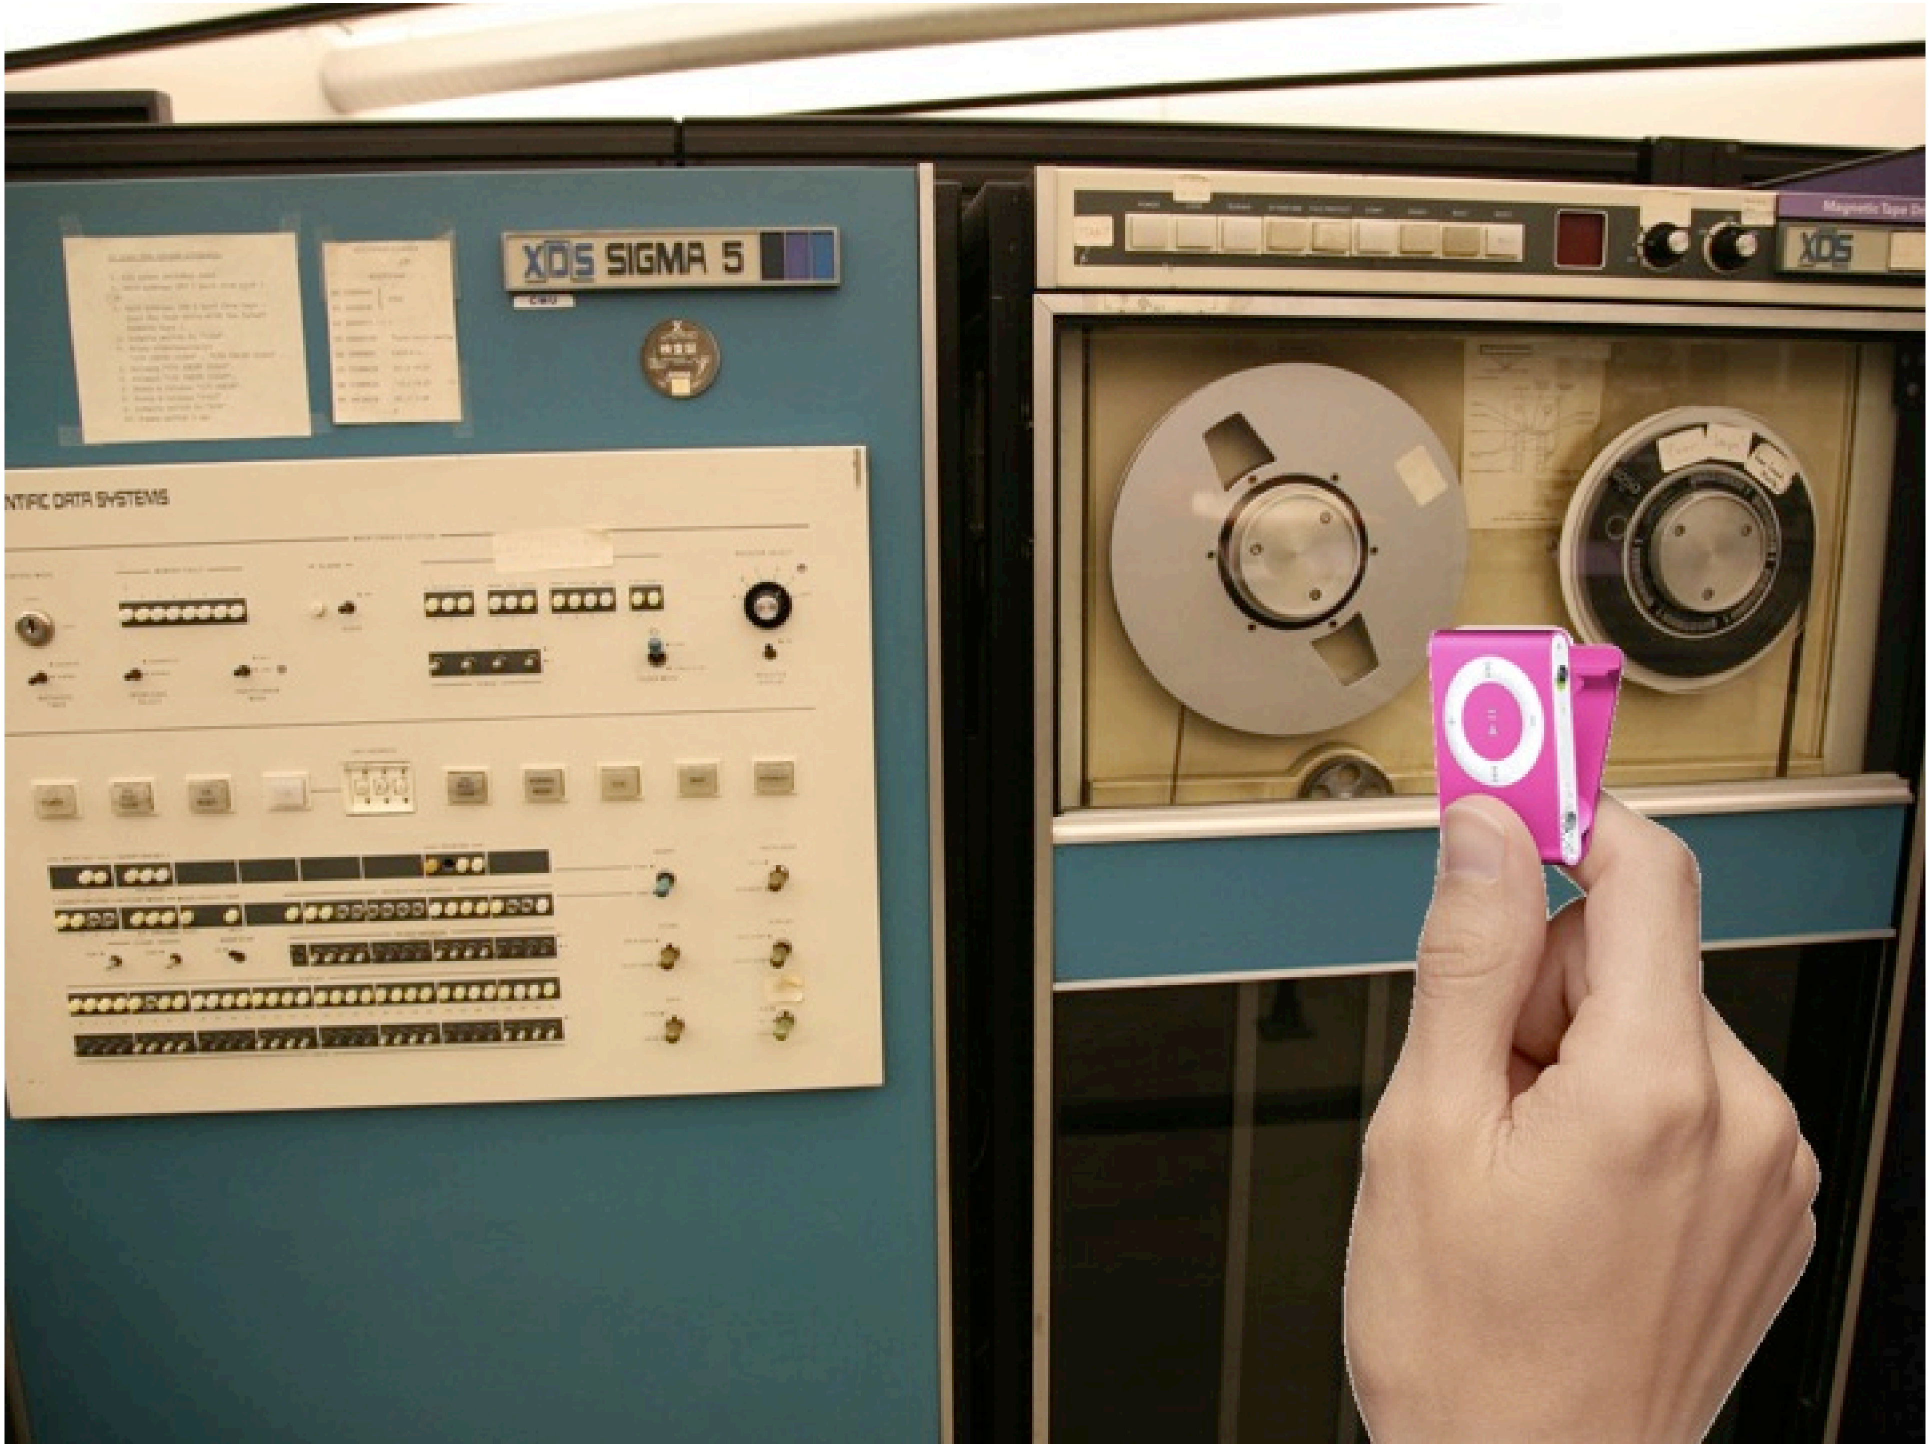
XDS SIGMA 5
XDS


Development of Multibeaming at GMRT
Goal: phased array sensitivity across full field of view

Australian SKA Pathfinder

- Focal Plane Array on 36 antennae
  - each dish: 30 x 1 square degree field of view
- Single Digital Backend (SDB) demonstration
  - concept proposed and pursued by Tim Cornell et al.
  - beam formation, correlation, and imaging on supercomputer
  - compute: 2 Pflops
  - transfer: 2 Tb/s in

Summary

- Modern backends combine hardware and software elements:
  - ☐ Hardware (FPGA) provides more grunt
    - enables large problems to be divided
    - increasing modularity and re-use
  - ☐ Software (CPU/GPU) provides more flexibilty
    - enables experimentation with more complicated algorithms
    - increasing performance
- Existing systems have great potential for discovery
- Future systems look challenging now, but …

SWINBURNE UNIVERSITY OF TECHNOLOGY

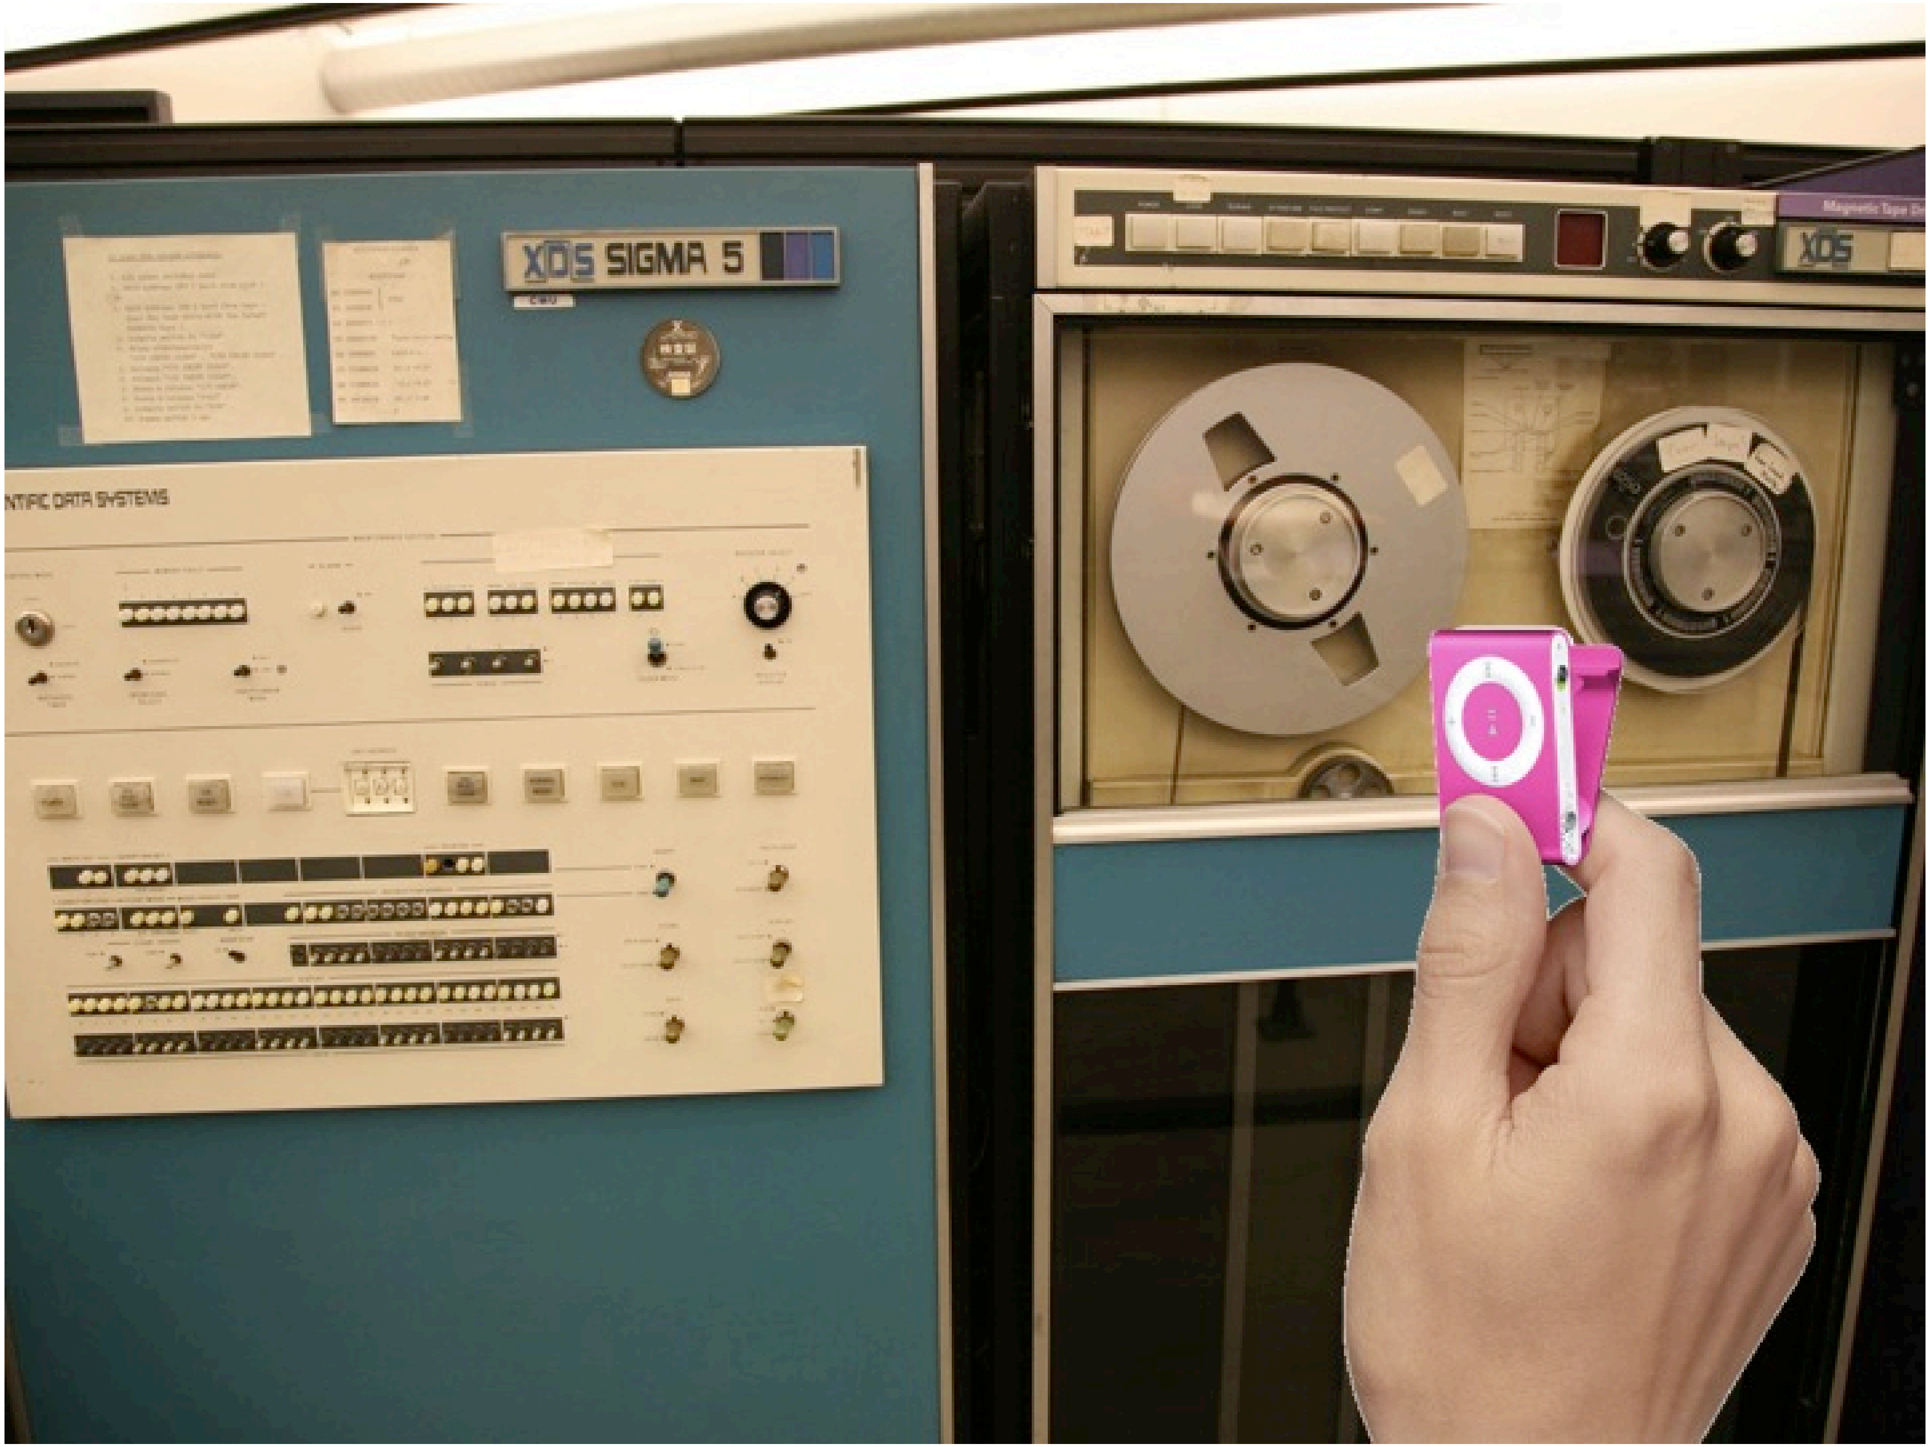
XDS SIGMA 5
CPU
XDS