Pulsar Timing – Backend Systems

Dr Tom Scragg

13/11/2024

DARA / HartRAO

Image credit: NASA

What are pulsars?

Image credit: NASA

The dying core of a supernova, the cooling heart of a deceased star sending a last fading signal out into interstellar space…

They are **not** stars that pulse. .
A pulsar is a neutron star where the magnetic pole is not aligned with the axis of rotation which therefore generates a beam of radiation that periodically aligns with our line of sight

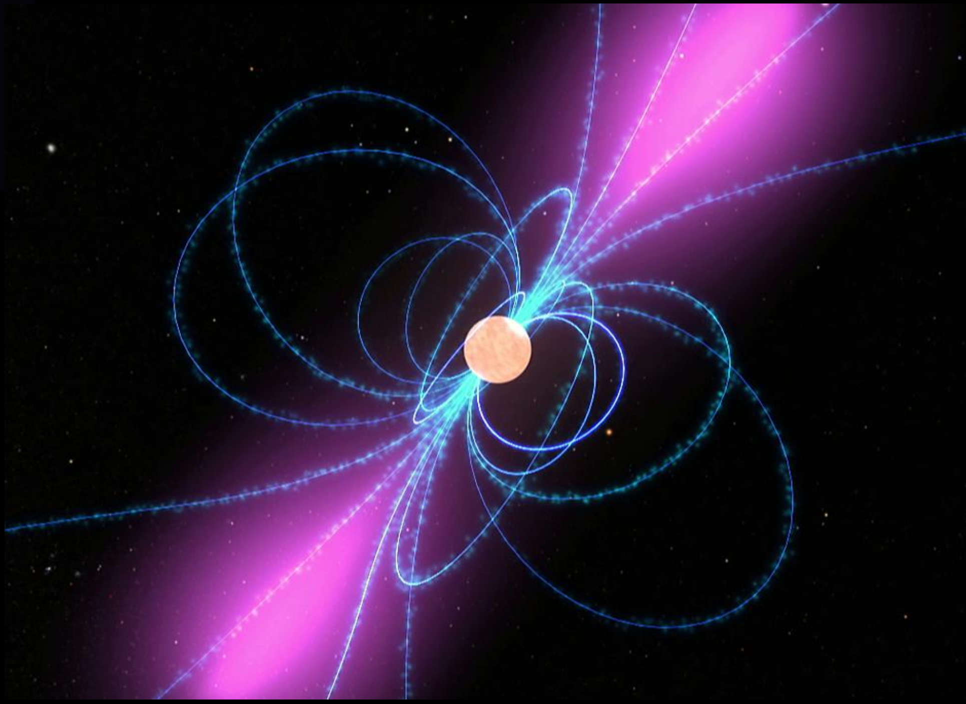

Image Credit: NASA/GSFC

Physics with Pulsars

Dispersion : Time of arrival at different frequencies

Frequency dispersion along the line of sight gives us information on the neutral hydrogen / electron number density of the ISM

By looking at the pulse profile

Changes in flux or phase or shape, we can derive structure in the emitting regions

Pulsar timing

Spin down rate, RRATs & Nulling emission mode, Glitches (neutron star interiors)

Pulsars in Binary Systems

Interactions between companions (Recycled pulsars Redbacks & Black Widows)

Millisecond pulsars

Post Keplerian Orbital mechanics, tests of General Relativity,

Constrain theories on gravity waves

Analogue Systems

Chart Recorder

- Simple and cheap
- Records are rolls of paper
- Sensitivity / resolution limited by speed of pen.

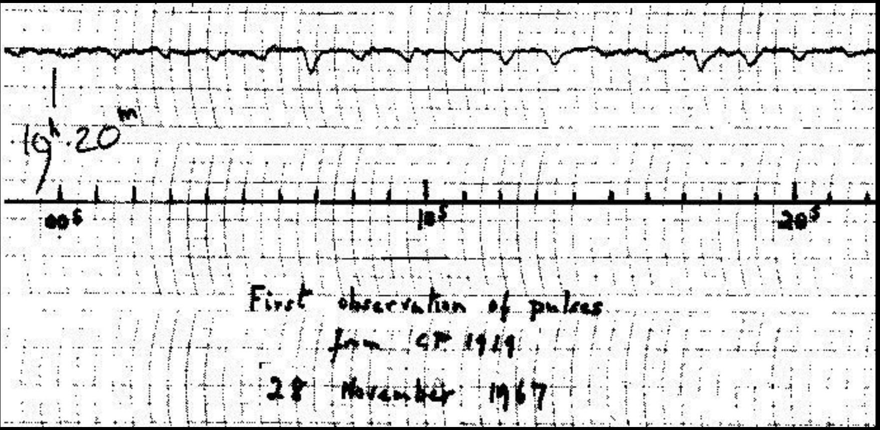

Image Credit: Adapted from Handbook of Pulsar Astronomy

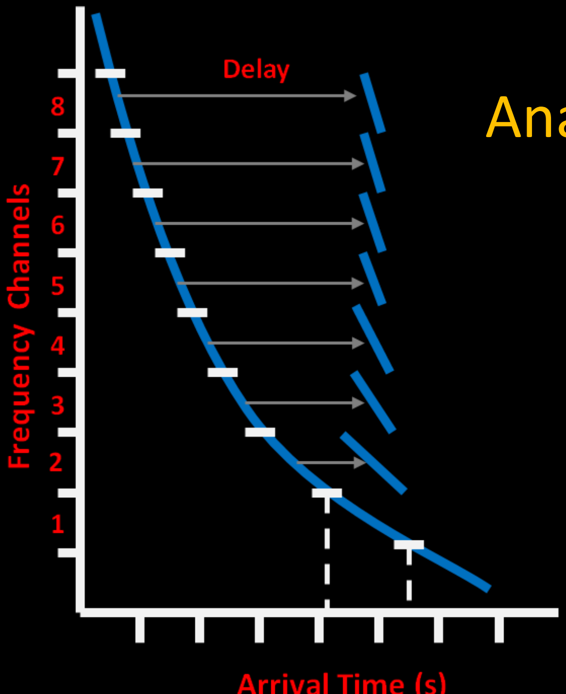

Analogue Filterbank (incoherent de-dispersion)

- Dispersion by the Interstellar medium ‘smears’ the signal in time and reduces the signal to noise ratio.
- Dedispersion reduces the above by delaying the signal by different times at different frequencies.
- Expensive / Custom Design
- Limited numbers of channels
- Tricky to setup

Digital Filterbanks

- Uses bespoke Application Specific Integrated Circuits and PFGA's for incoherent dedispersion on many frequency channels.
  - Dedicated and Fast
  - Configurable (but not flexible)
  - New algorithm = new ASIC or FPGA
  - Approximates the cost of a new system
- Expensive and Time consuming to design

General Purpose Computers

- As CPU compute power increases (Gflops) then ‘low cost’ general purpose computers can be used.
  - Flexible and programmable
  - Cheaper (off the shelf components)
  - Upgradable
  - Parallel processing
- Limitations
  - Input / Output speed
  - Memory capacity
  - Memory bus limited
- Use Beowulf Clusters to process wider bandwidths / higher data rates
- Implement coherent dedispersion & folding

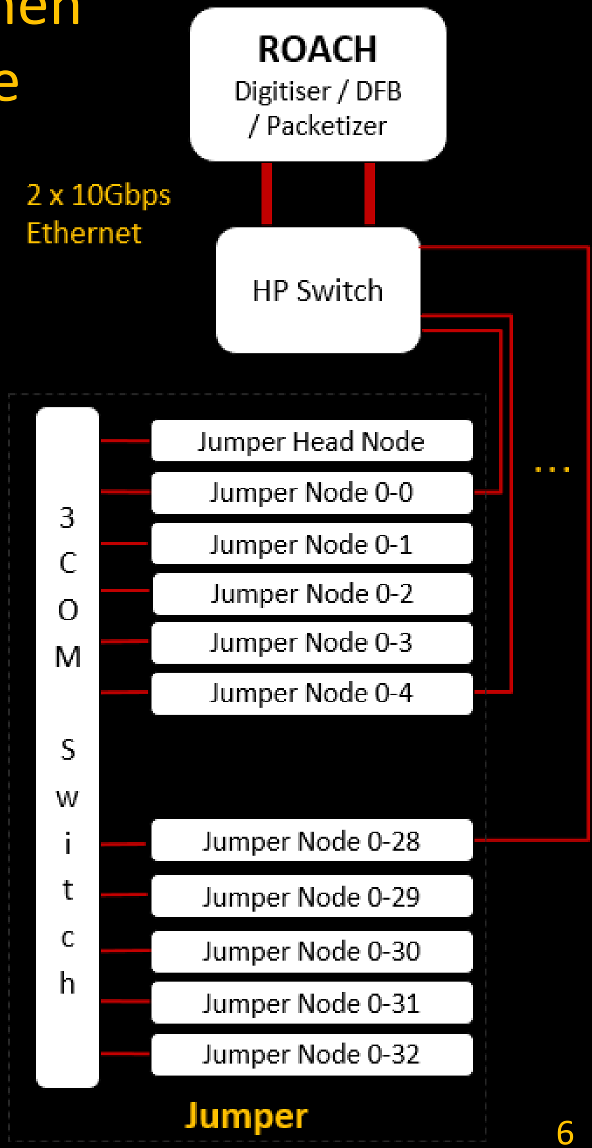

Coherent Dedispersion

- Coherent Dedispersion
  - Use a chirp filter to reverse the effect of the Interstellar Medium on the pulsar signal.
  - Gives a much better signal to noise and pulse shape resolution
  - Use a FFT to carry out dedispersion and filterbank synthesis
  - The processing load is proportional to the Fourier transform length.

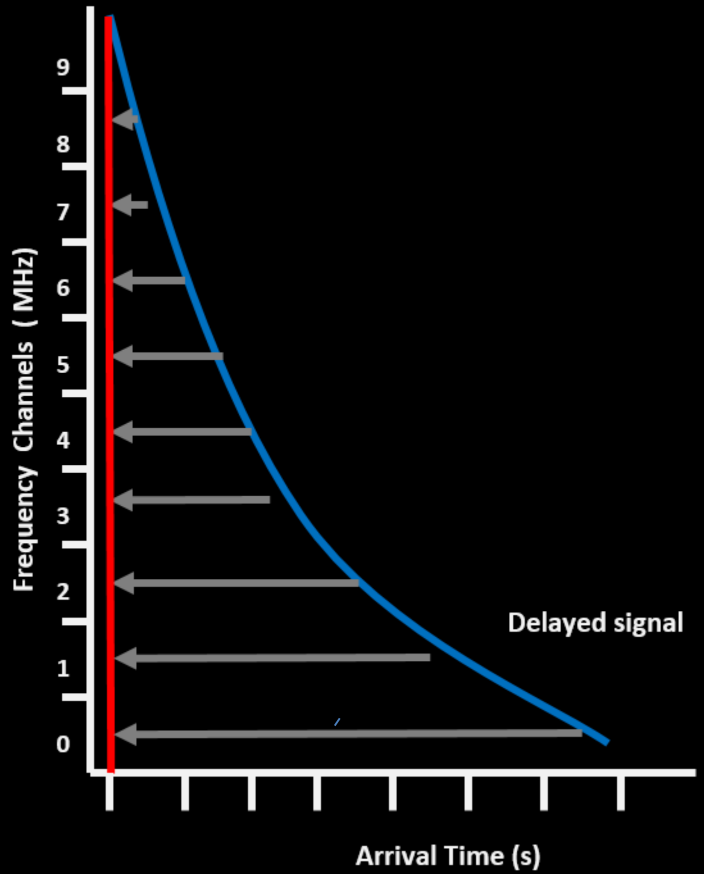

- Undertake processing in real time
  - Alternatively store the input data and process later
  - Accumulating at 1.536 GBytes per second*

Dedispersion

Incoherently dedispersed | **Coherently dedispersed**

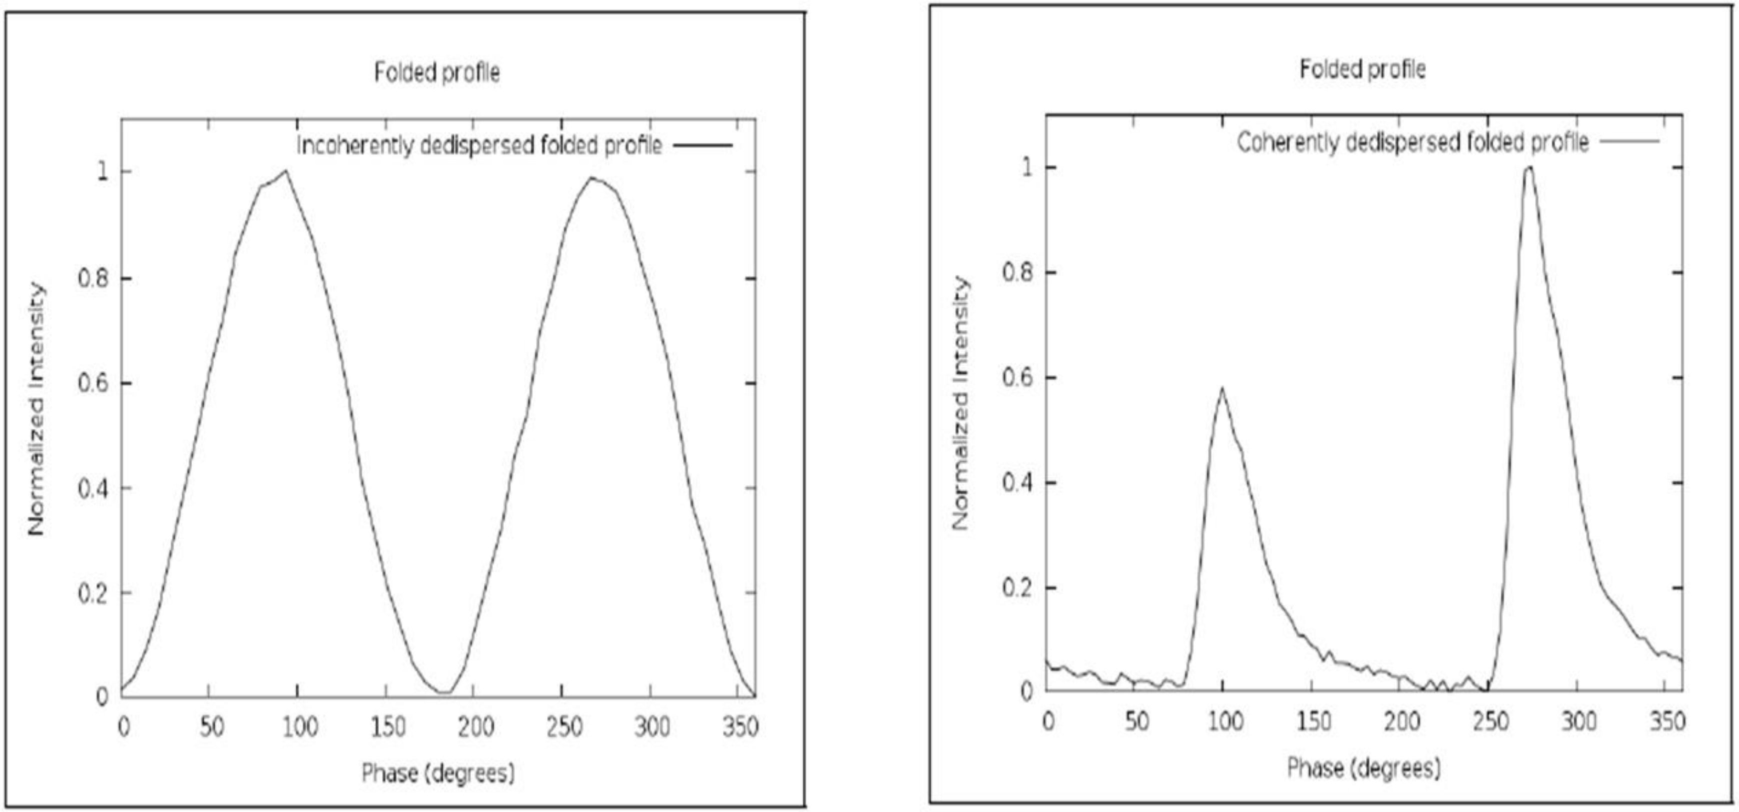

(a) Incoherently dedispersed and normalized folded profile for B1937+21.

(b) Coherently dedispersed and normalized folded profile for B1937+21.

Image Credit: Adapted from Handbook of Pulsar Astronomy

Scrunching and Folding

- The noise within a radio signal is not correlated with respect to time or frequency.
- Scrunching in frequency
  - Filter banks separate the incoming signal into frequency channels
  - After dedispersion the signal is correlated wrt to time but the noise is not
  - Combining the separate frequency channels provides improved S/N.
- Folding in time:
  - We observe the pulsar for a much longer time than its period of rotation.
  - Adding multiple periods together reduces the noise and provides signal gain (S/N)
  - *And a stable pulse profile*

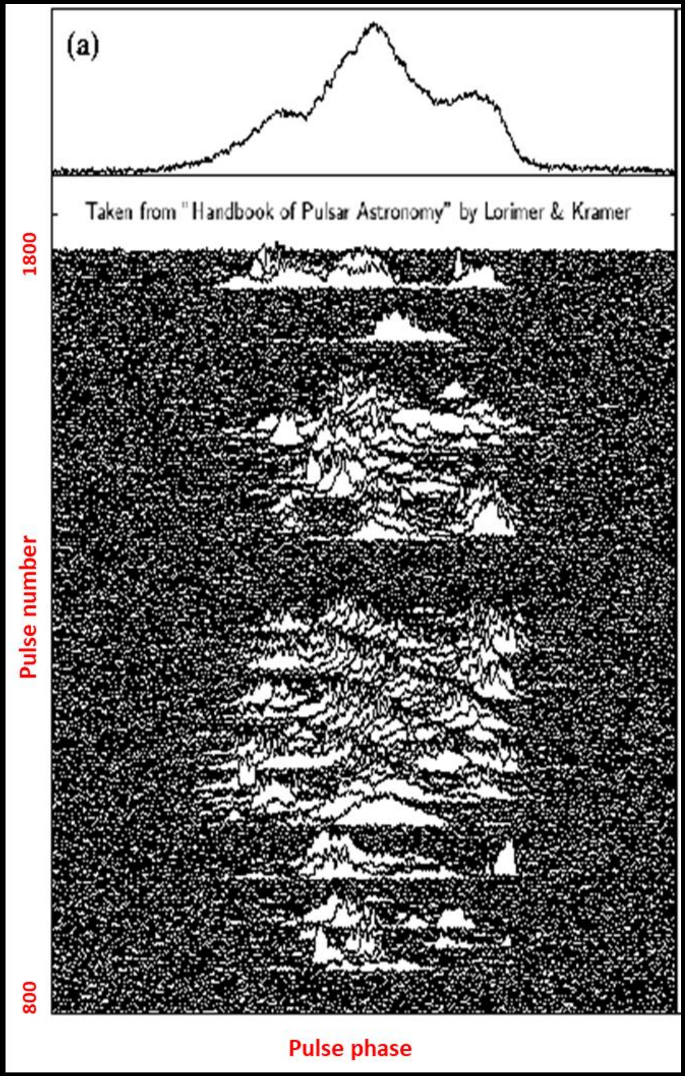

Ever Increasing Data Processing

Data Rates

Events duration ~ 100's ns, high bandwidth and fast sampling to get that time resolution

512MHz sampled at 1 GHz = 1GByte of data per second

Volume of Data

Normal pulsar: 1 s period ~ 1 GByte of data to dedisperse per pulse

Millisecond pulsar: 10ms period ~ 10MBytes of data to dedisperse per pulse

High Dispersion Measure Pulsars

A pulse may be spread over many periods

Processing the data in Real Time

Each channel ~ 2 x 32 Mbytes per second

Each channel processed by a separate node

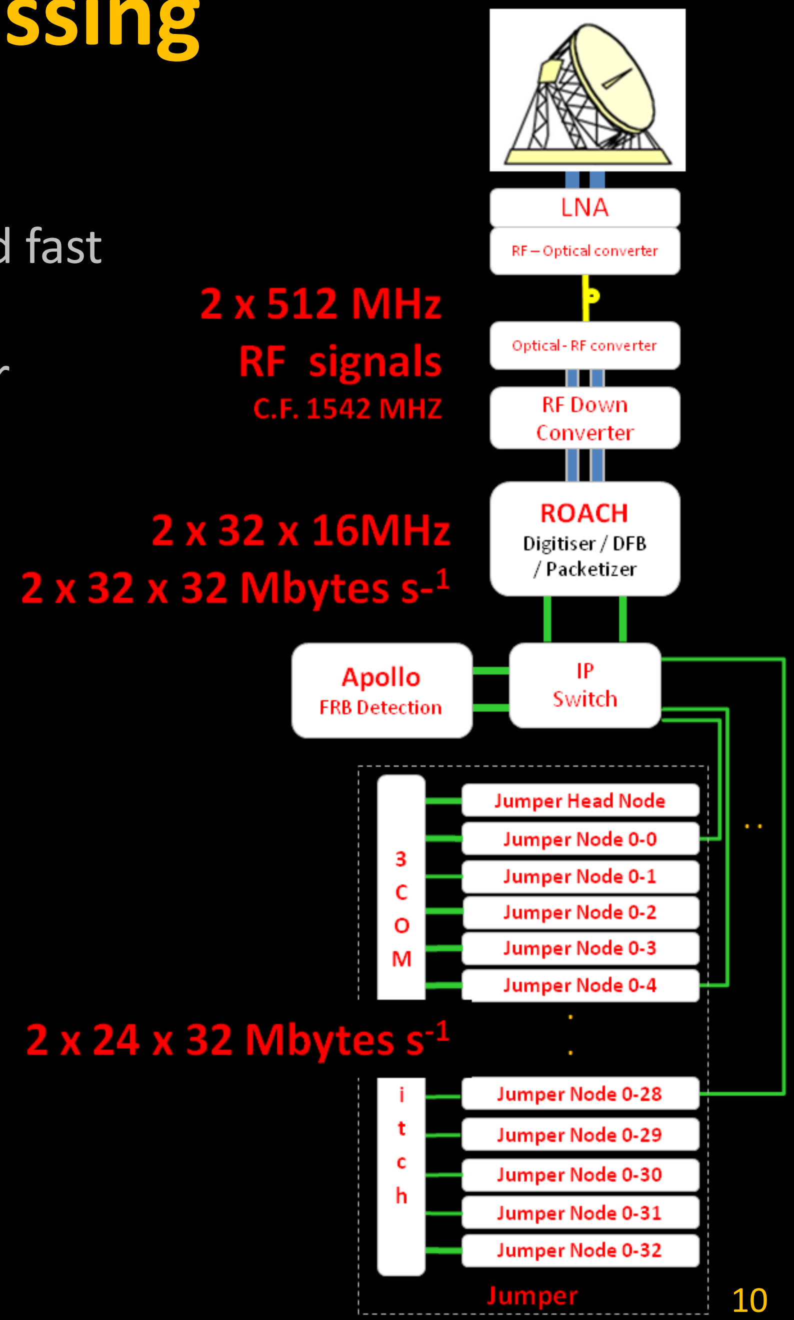

GPU Based Processing

- Massively parallel processing, for a correctly designed program GPU>>>CPU
  - 1000 + ALU per card
  - 12 Gbytes memory per card
  - Teraflops of processing capacity per card
- Independent computational operations in each core
  - reduced synchronisation and signalling requirements
  - faster bulk transfer of data rather than single memory read/write operations
- Contention for shared resources (intermediate cache memory, system transfer buses) means overall performance x ~ 1.6.
- Limitations
  - Do Not Branch ! – requires slow main memory accesses
  - As before memory and bandwidth limitations

Hebe: Pulsar Timing System

GRAO and JBO

Hybrid Cluster of compute elements

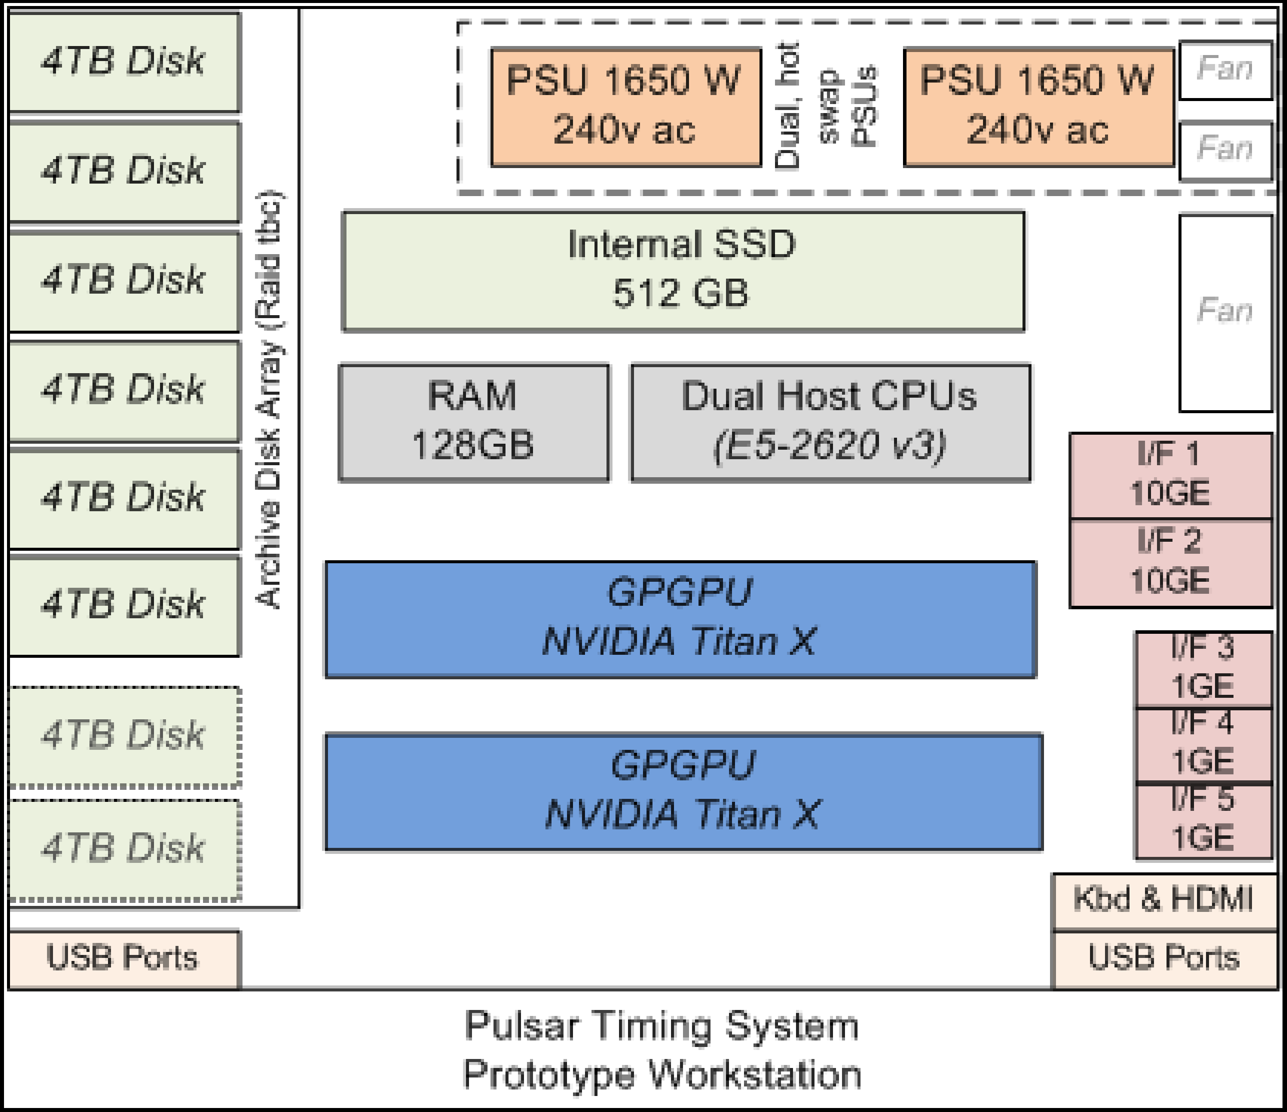

Computing Architecture

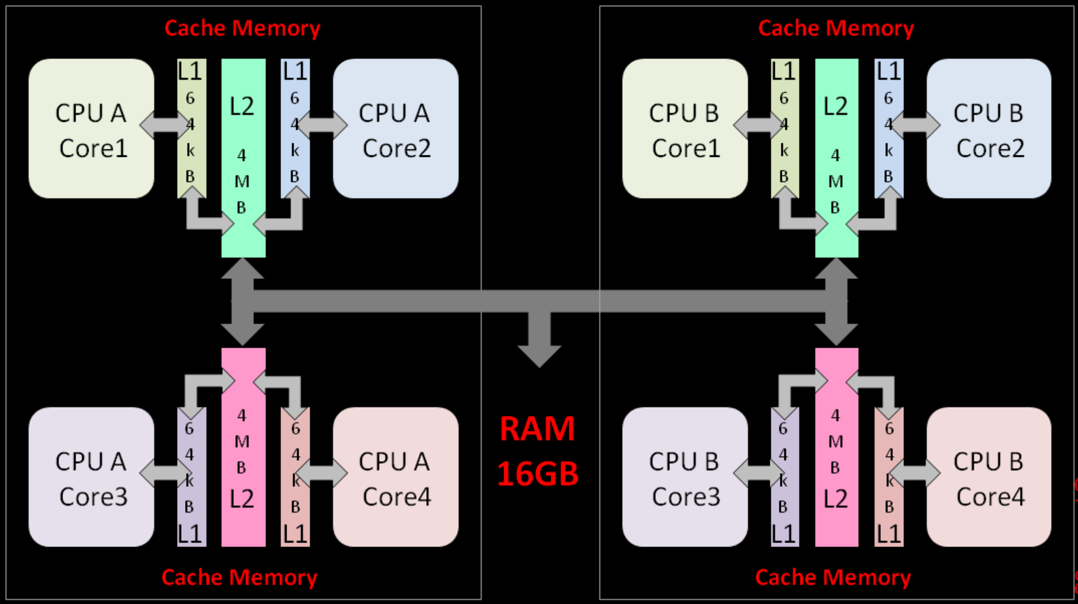

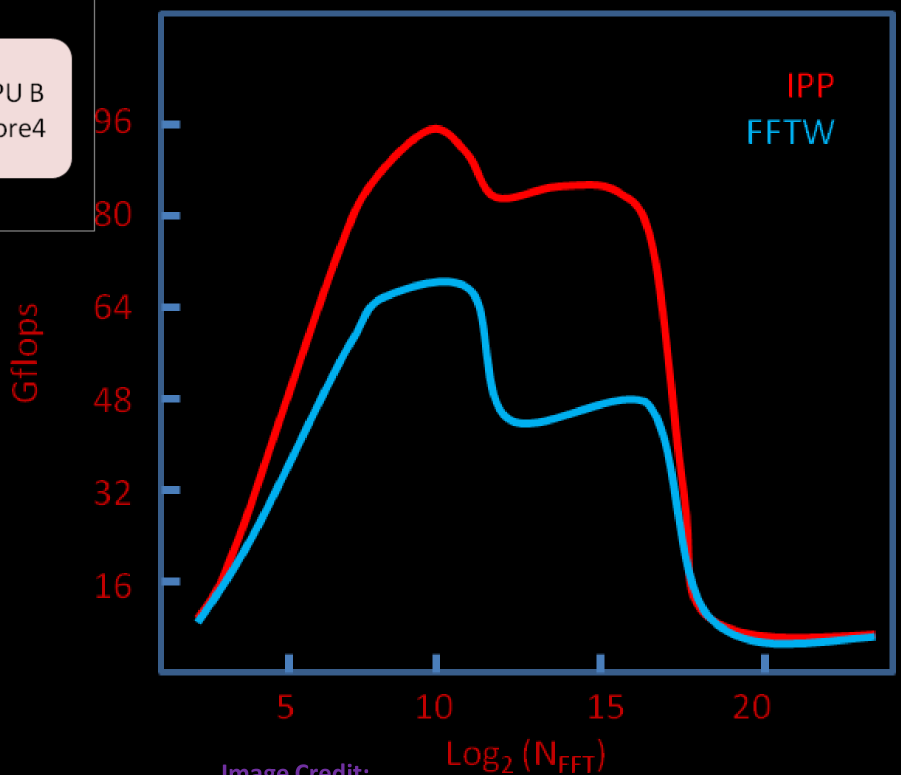

Processing throughput can collapse

Software and hardware important

N_{FFT} ~ array size for FFT algorithm

Based on Dual Quad Core Xeon processor workstation

Image Credit:
Adapted from van Straten and Bailes 2010

GPU Architecture

Nvidia Fermi GPU architecture

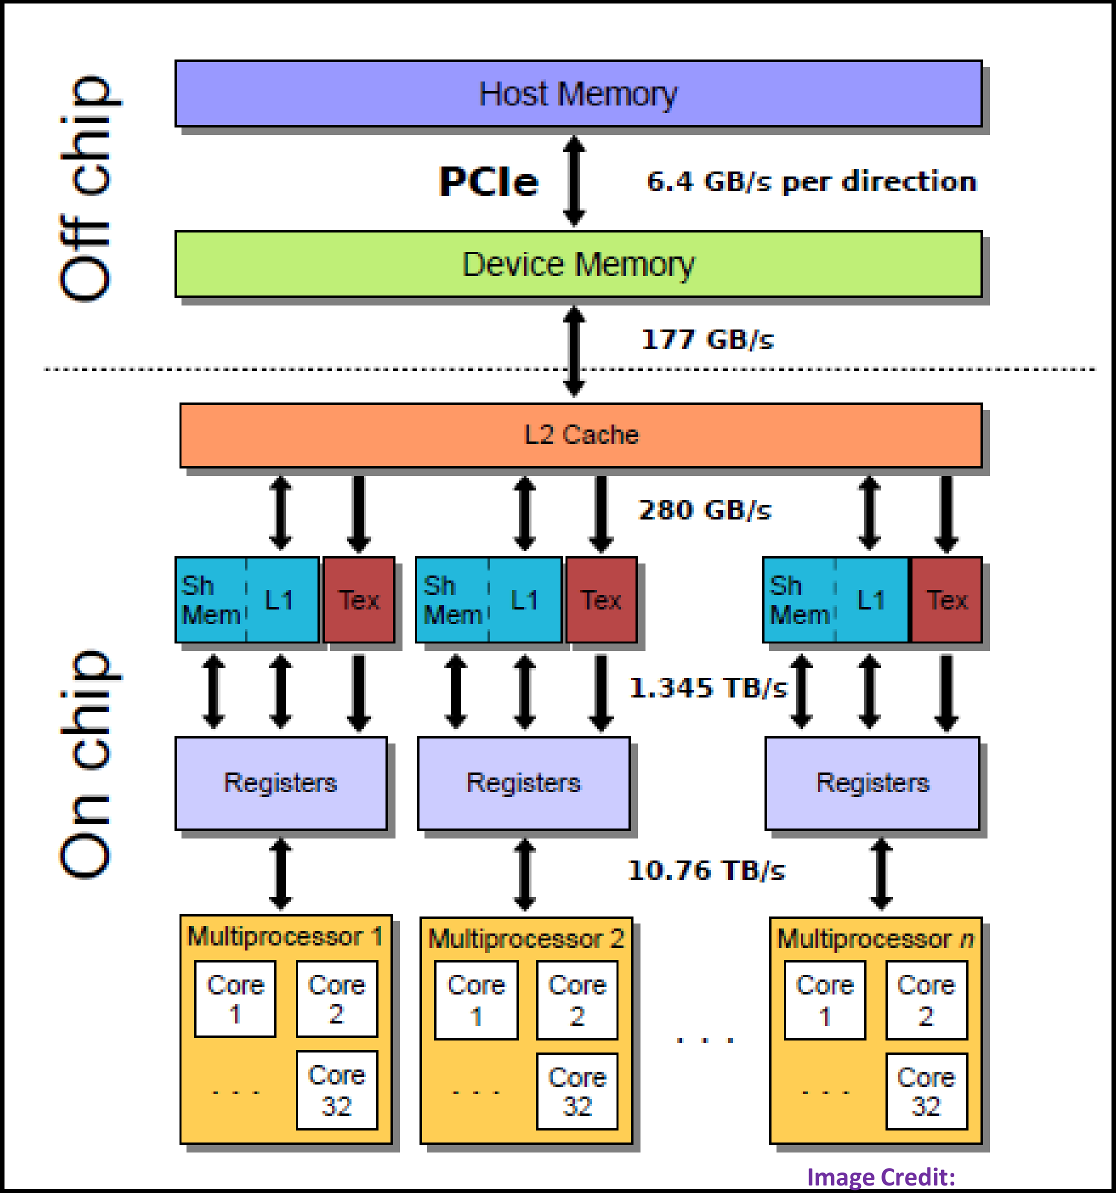

Image Credit: adapted from Clarck, LaPlante & Greeenhill 2011

Not All GPUs are equal

Two Types of GPU devices

- Gaming : more powerful for compute intensive operations
- Scientific : better I/O capabilities
- Five fold price difference

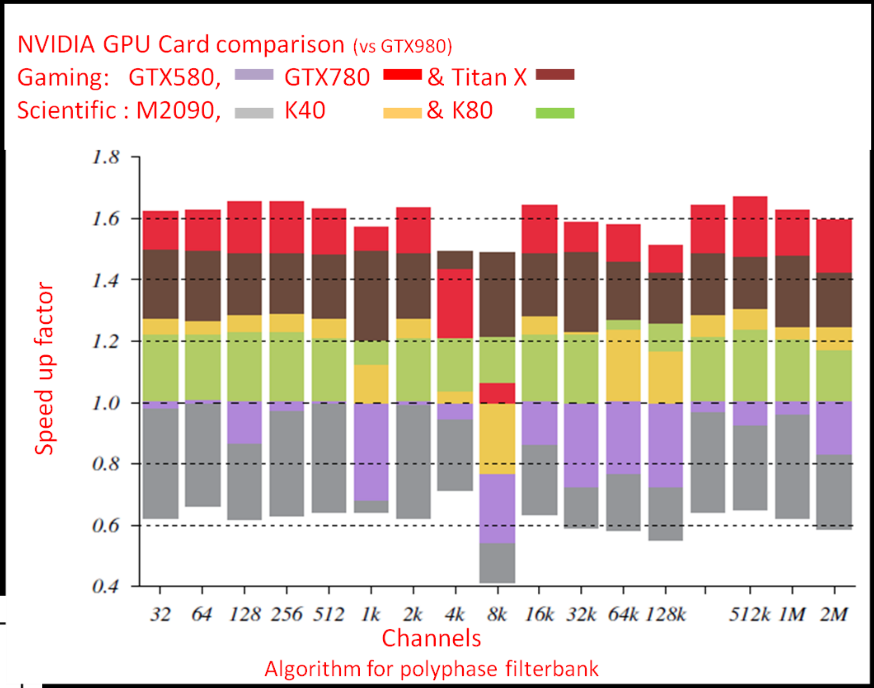

Image Credit:
adapted from Adamek, Novotny & Armour 2015

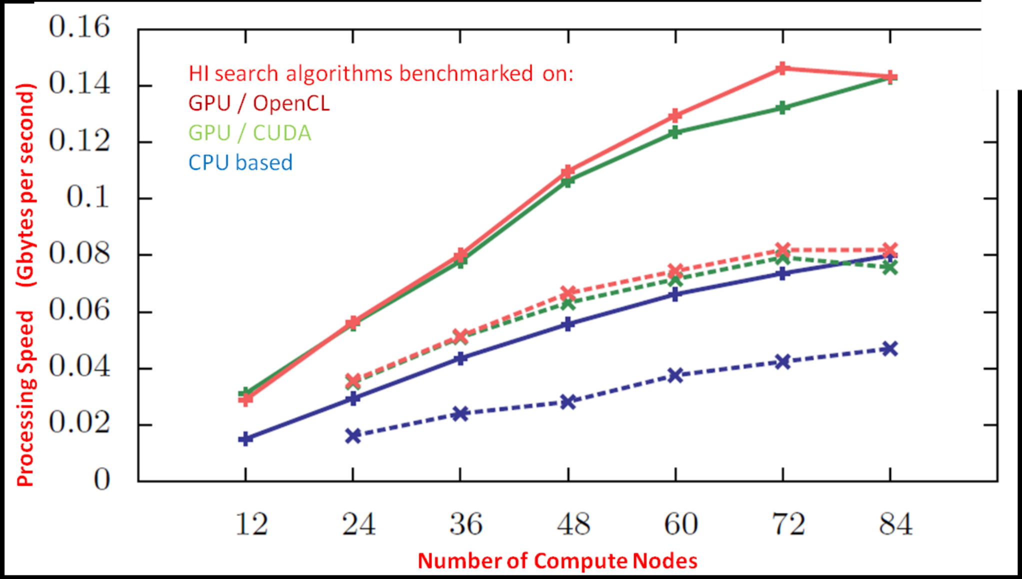

Impact of Development Environment:

- CUDA : NVIDIA Proprietary
- OpenCL : Open standard

Image Credit: Adapted from Westerlund & Harris 2015

Future Demands

- Multi-beam / Multiple receivers
  - More sky coverage
  - More data to process / more GFLOPS required
  - Lower pointing accuracy / rotation of beam on the sky
- Phased Array Feeds
  - Many more beams
  - Beams move on the sky as the telescope tracks central source
- Real time beamforming very difficult
  - Use pre computed weightings?
- Ultra wide bandwidths
- Signal processing of data from the individual beams.

Thank you