ARTEMIS and GPU based real-time signal detection.

W. Armour, A. Karastergiou, C. Williams and M. Giles.
University of Oxford.

13th November 2013

Talk: Part I…

First: An introduction to ARTEMIS and our current results.

Second: ARTEMIS – what's under the hood? (GPU de-dispersion using Astro-Accelerate).

What is ARTEMIS…

Real time detection system to detect RRATs and Lorimer bursts as they happen.

Built around affordable and scalable hardware with accelerated processing provided by NVIDIA GPUs.

Buffered voltage and Stokes data allows for extensive follow-up analysis.

What is ARTEMIS…

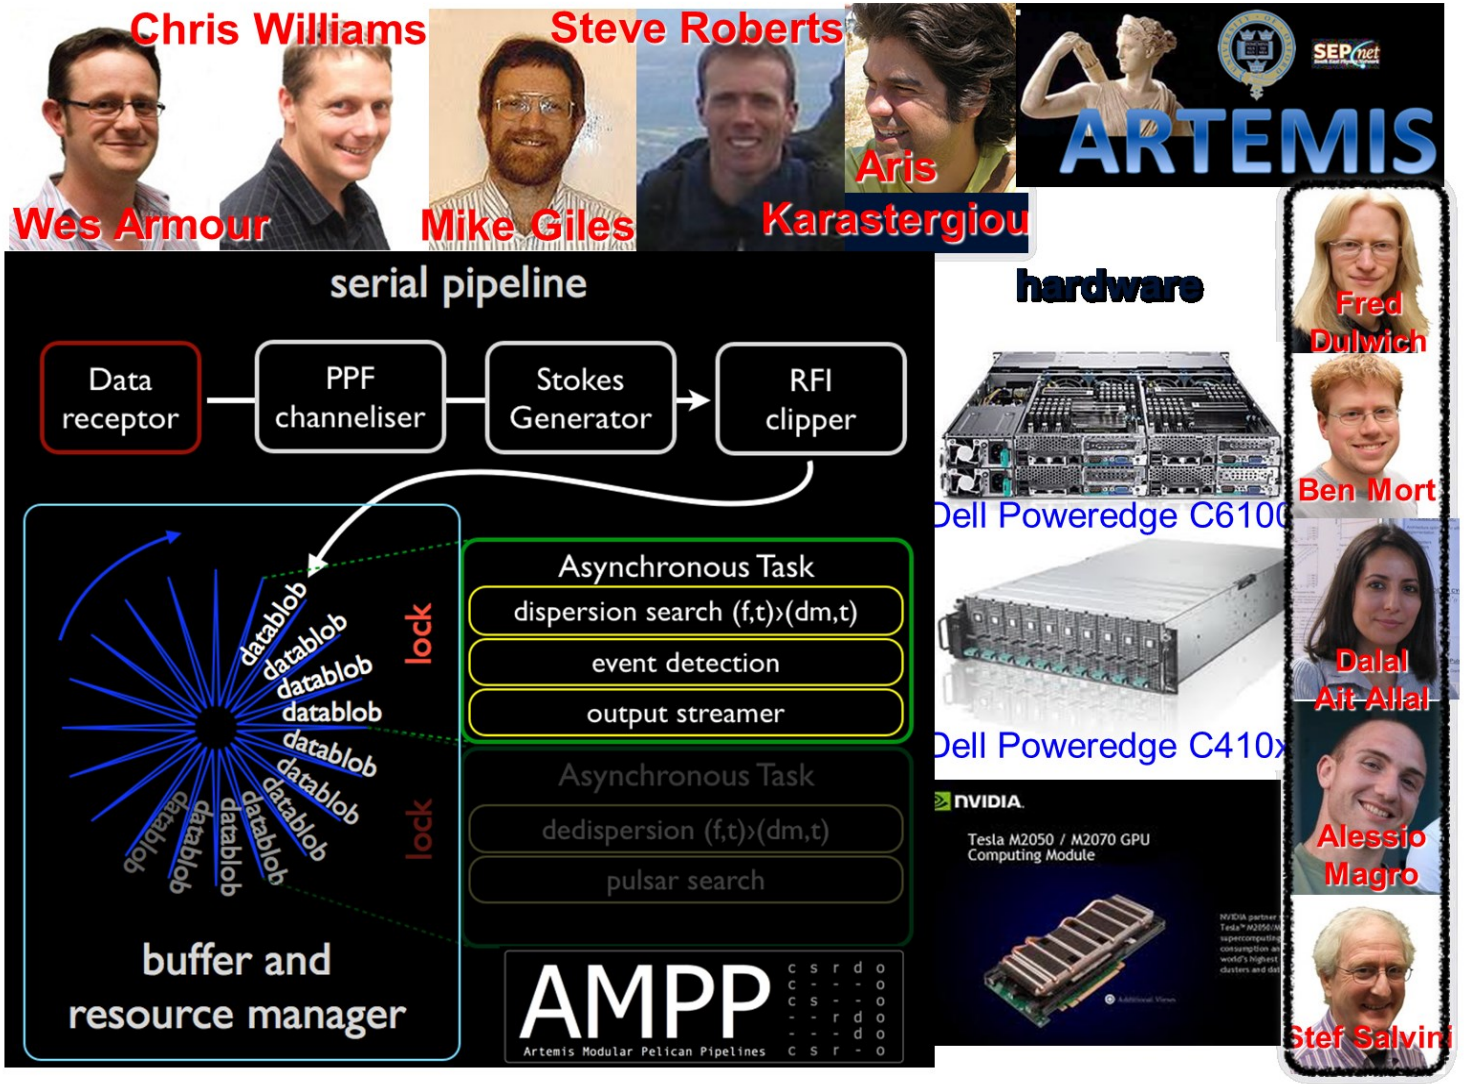

ARTEMIS and LOFAR international stations…

Large collecting area, high sensitivity

96 coherently added dipoles x 2 pol

Very fast sampling rate: 5.12μs voltage data

Large fields of view: 30 sq.deg in 8 x 6MHz beams (at 150 MHz)

Large amounts of observing available time

ARTEMIS and RFI mitigation…

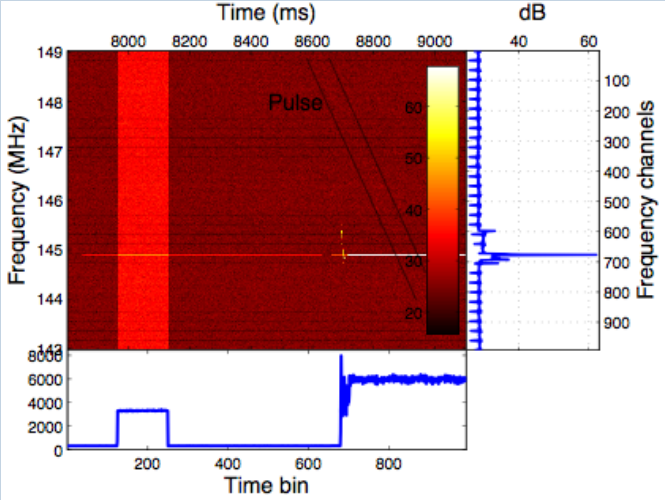

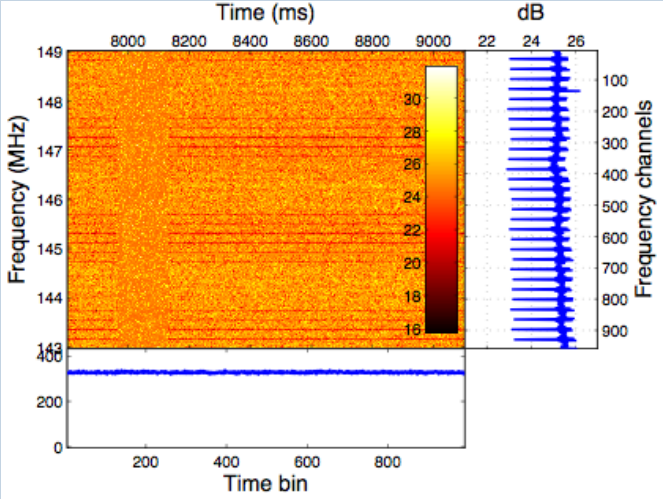

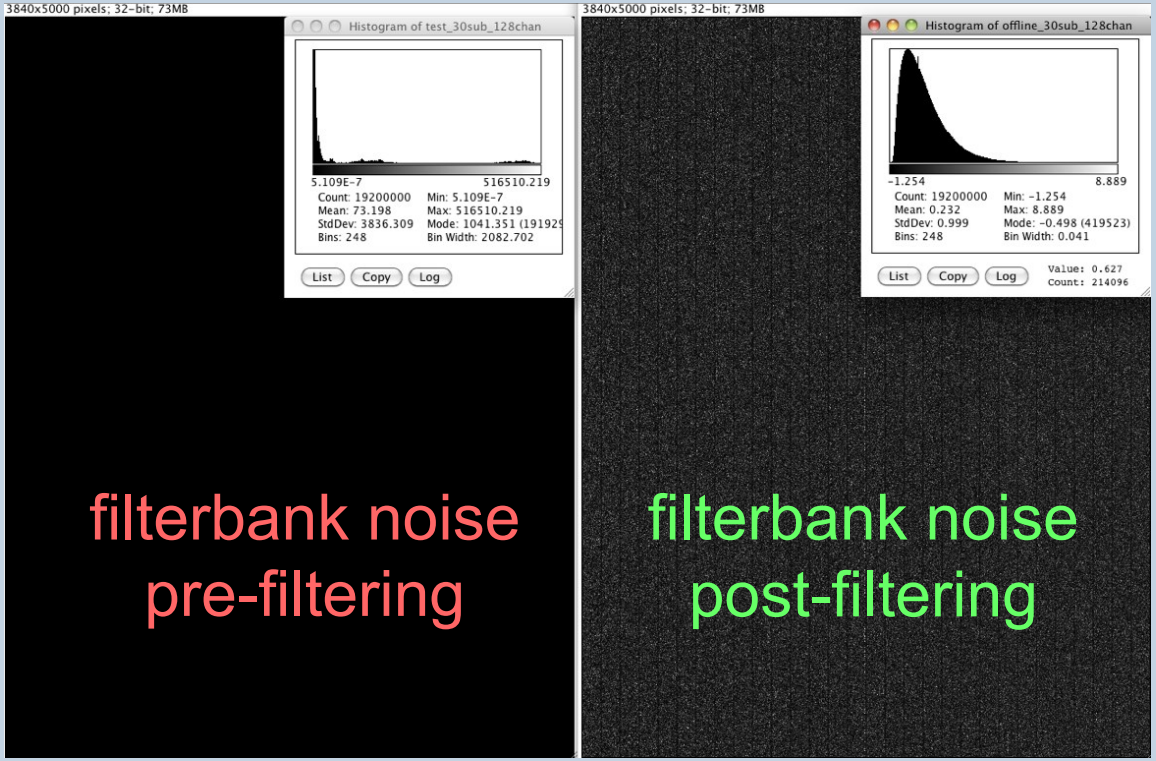

Dynamic noise modeling

=

Effective RFI removal

OXFORD e-Research CENTRE

ARTEMIS and Many-Core Acceleration…

GPU de-dispersion test:

1500 measures searched in real-time, only pulses from B0329+54 detected, no spurious signals

(Will talk more about this later)

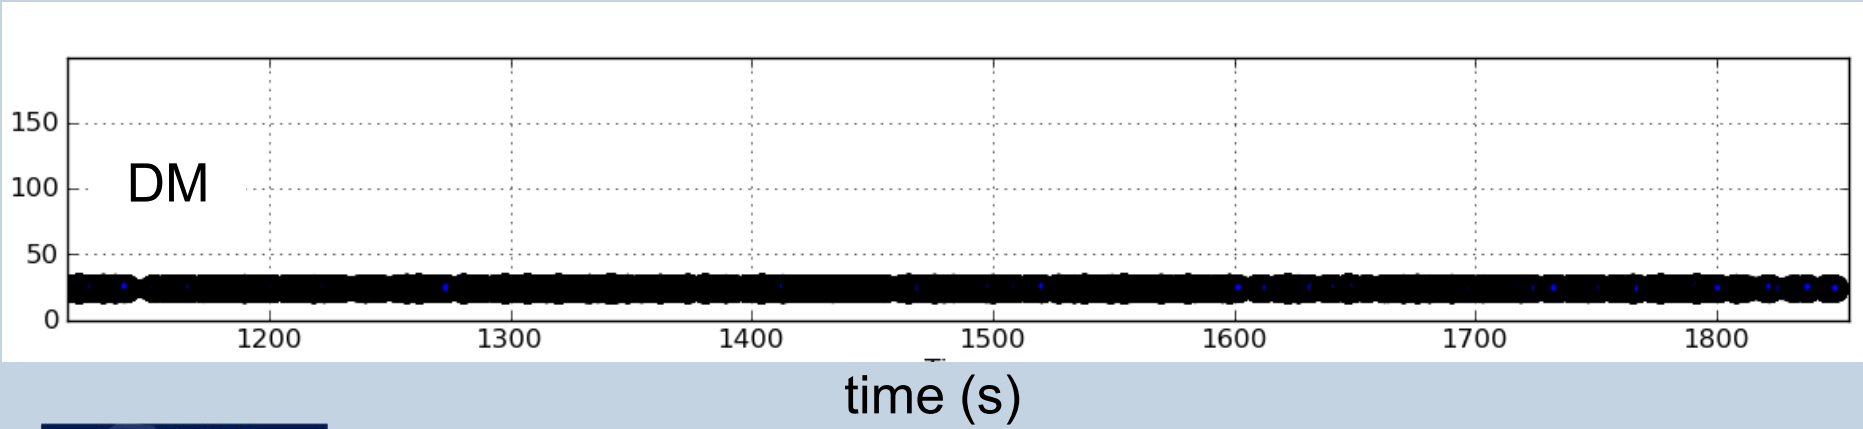

ARTEMIS: Initial Results…

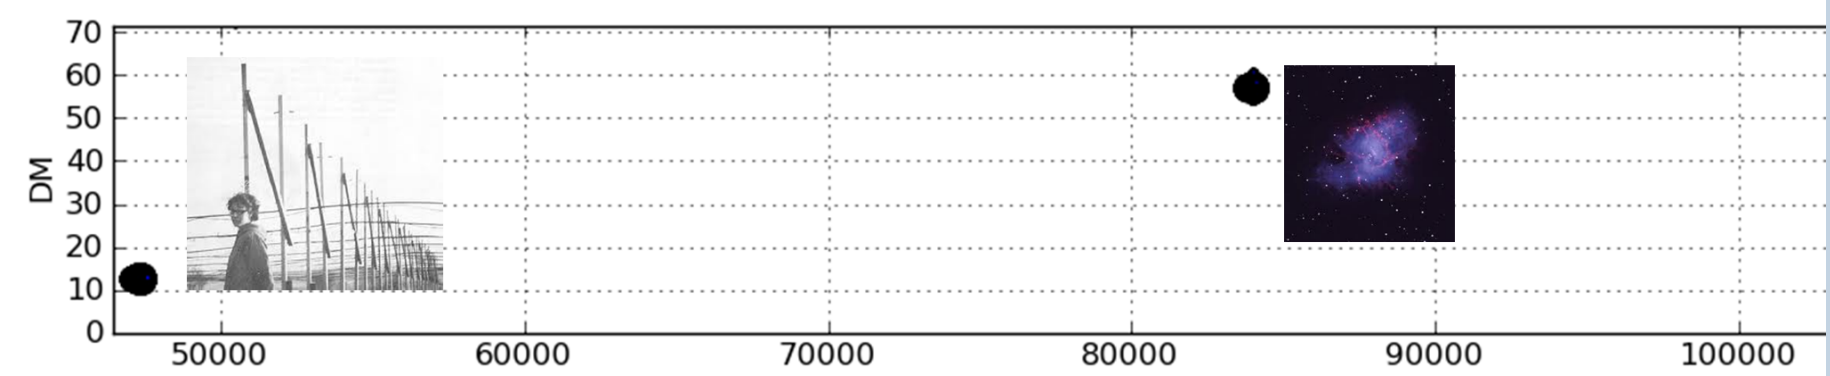

26 December 2011
drift scan with beam at Dec=+21º

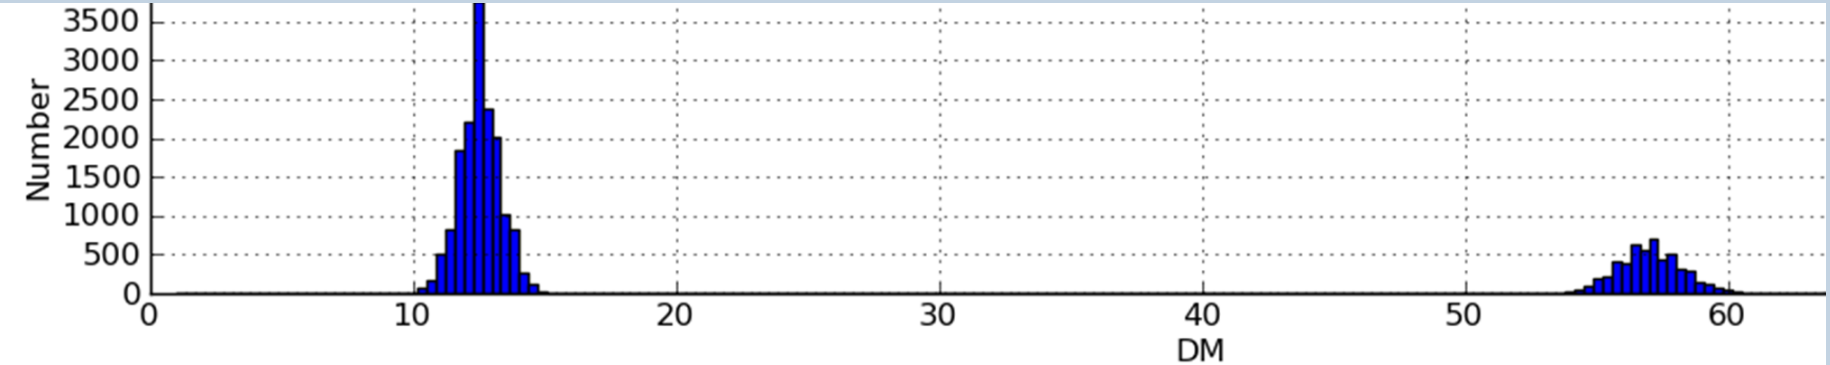

OXFORD e-Research CENTRE

ARTEMIS: Survey I…

- 1000 h of observations
- 8 beams on Greenwich meridian
- 143 - 149 MHz
- ~2000 x 3.05 kHz channels
- ~40 Jy (at 8σ) over 2ms integrations
- DM_{max} = 150 pc / cm^3
- Sampling local Galactic population (really?)

ARTEMIS: Survey I, Results…

For 40 Jy (8σ) sources in 2ms integrations:

$R < 1/1000h/30deg^2 \approx 30$ sources / sky / day

$\approx 3x10^{-5}$ sources / deg^2 / hour

Consistent with Lorimer et.al. 2013 and Hassall et.al. 2013

ARTEMIS: Current Status…

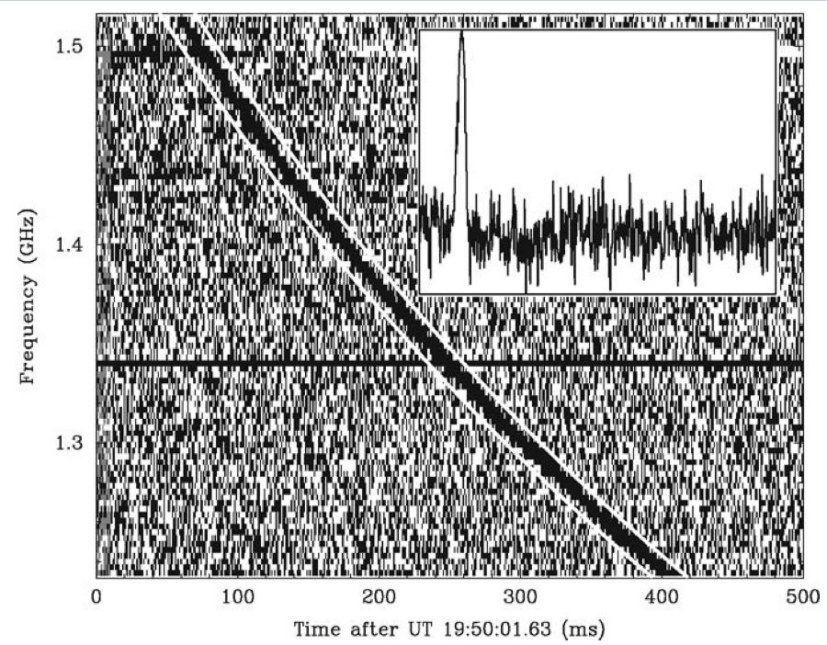

Lorimer 2007

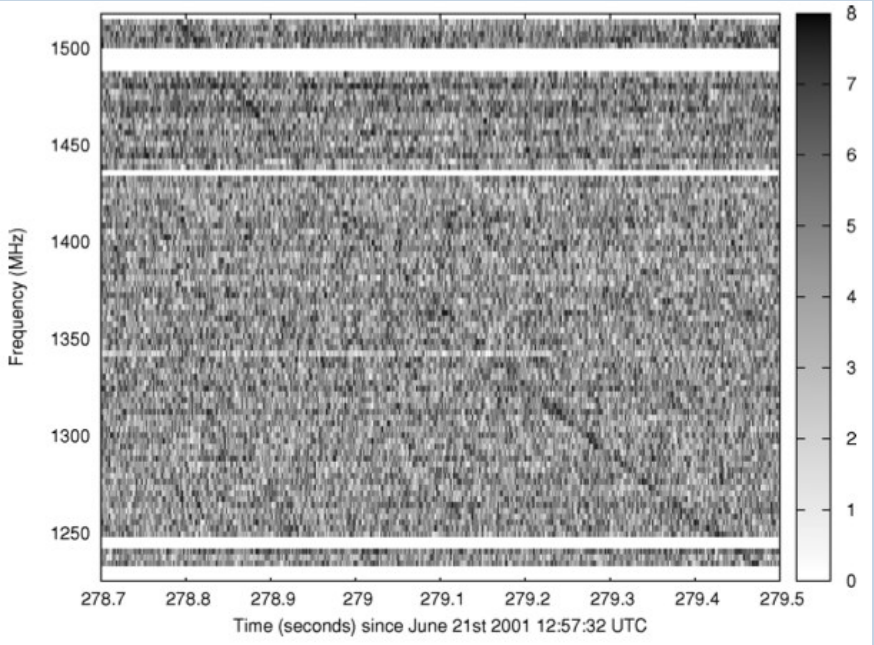

Keane 2012

ARTEMIS has no confirmed detections yet...

ARTEMIS: On going and Future Work - Survey II…

- 8 beams on Greenwich meridian
- 143 - 149 MHz
- ~4000 x 1.5 kHz channels
- ~650 microsecond resolution
- DM_{max} = 320 pc / cm^3

Talk: Part II…

First: An introduction to ARTEMIS and our current results.

Second: ARTEMIS – what's under the hood? (GPU de-dispersion using Astro-Accelerate).

Astro-Accelerate: Aims…

- To produce a many-core accelerated library for pulsar/FRB signal detection and data processing.
- Support multiple architectures such as GPUs, CPUs and Xeon Phi. With the ultimate aim of being hardware agnostic without a sacrifice in performance.
- Auto-optimizing for a given telescope/survey and computer configuration.
- For the library to be able to be called from existing codes such as AMPP(ATREMIS), sigproc or presto - code agnostic.

Features 1: Time binning…

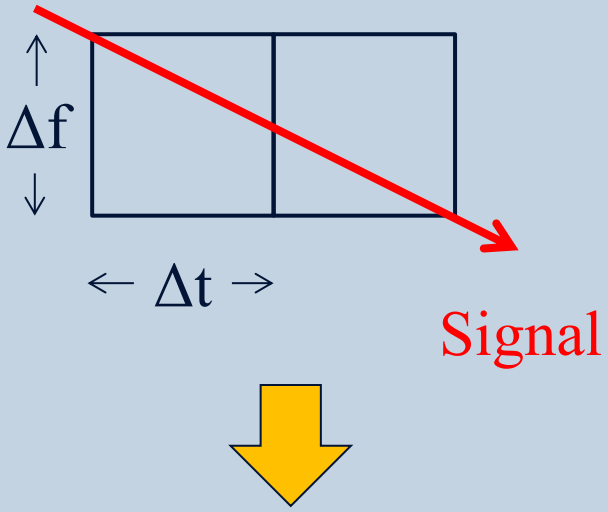

Has the effect of reducing the amount threads that are needed to process a region of (DM,t) space.

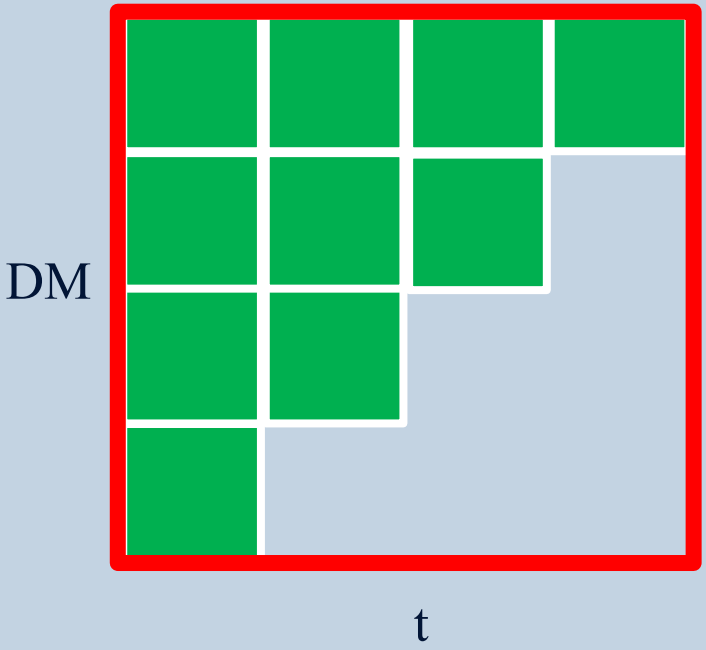

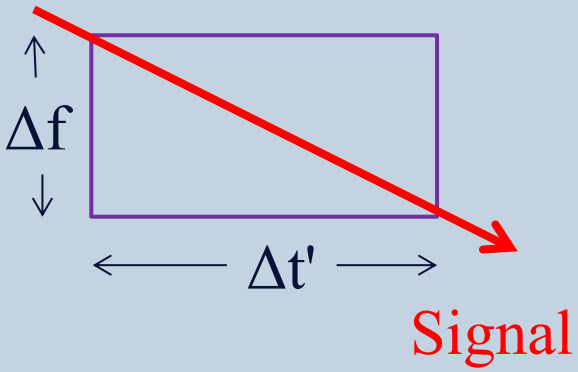

Utilizes the CPU and GPU at the same time (analyze previous de-dispersed data or bin on CPU).

Features 2: Code Execution Path

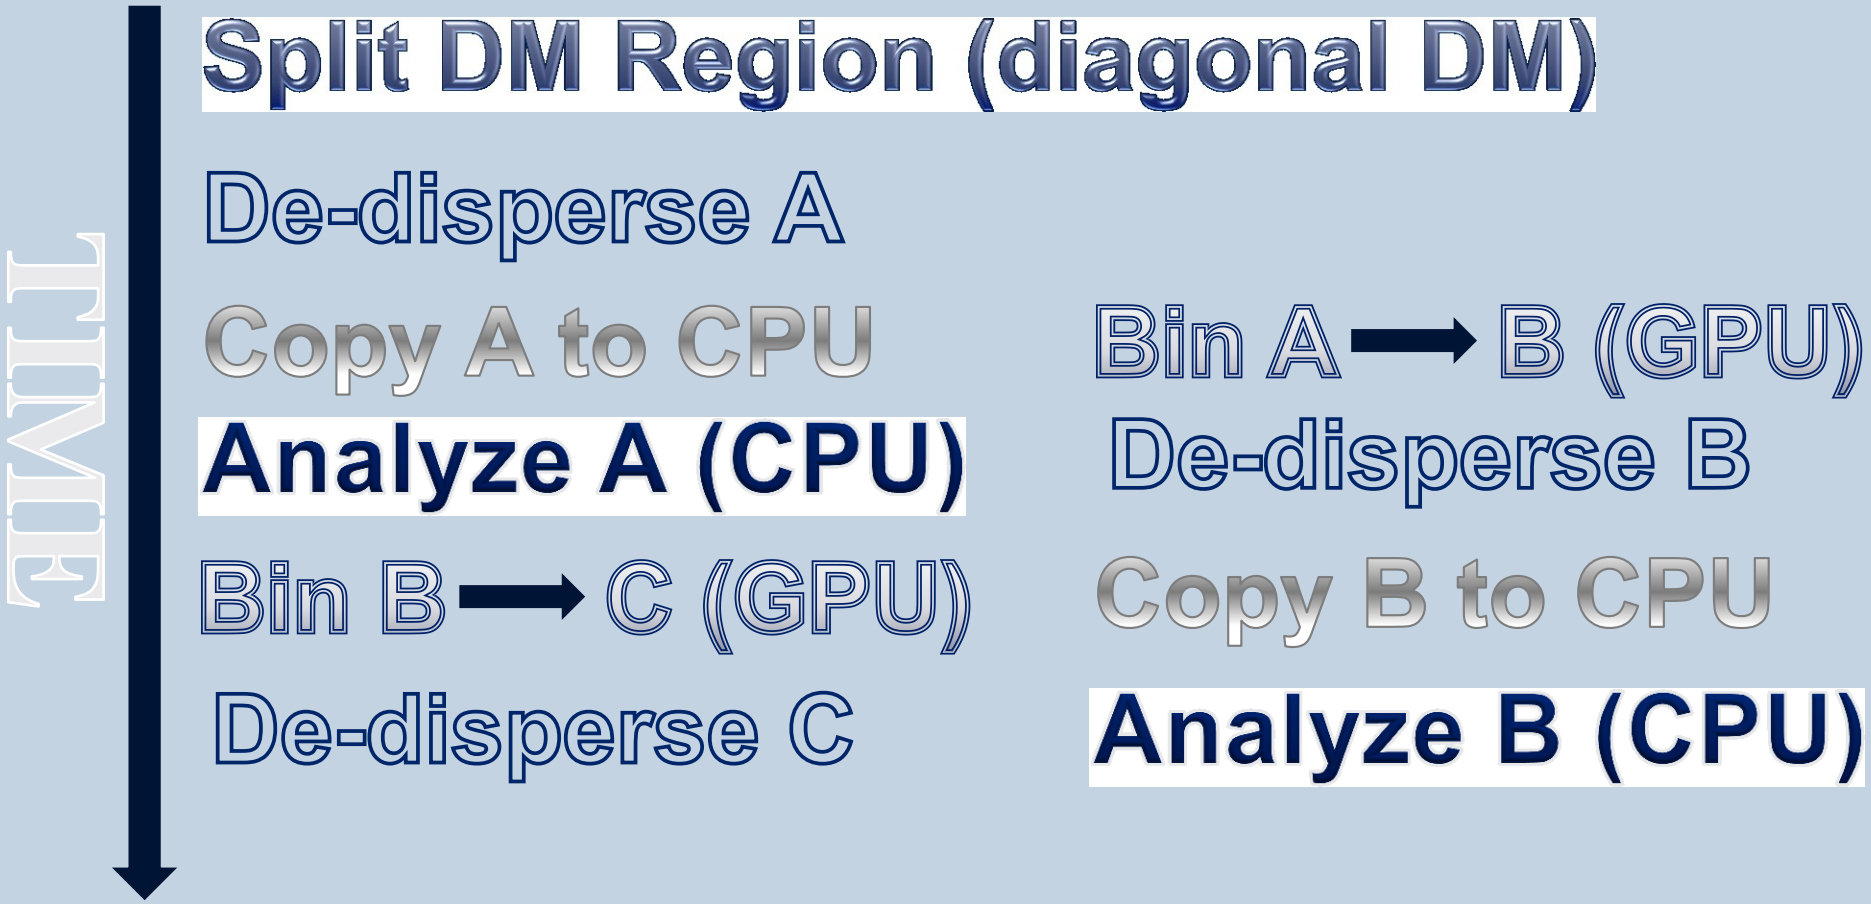

Features 2: Code Execution Path

Split data and DM regions to be processed:

Allows for Multi-core (and Vector) CPU usage along with PCIe and GPU usage at the same time = High system utilisation

Features 3: Kernels…

Different GPU kernels for different telescope/survey/hardware configurations…

- L1 Cache – most generic (slowest)
- Texture Cache – Kepler only (fast)
- Intrinsics – K20 only (faster)
- Shared memory – Fermi and Kepler (fastest, best used with time binning)

Features 4: Processing…

- Can search any power of frequency – **not just $1/f^2$**
- Can measure **negative** DMs (slower kernel only at the moment)
- Can switch analysis on/off
- Can dump (dm,t) output to disk

De-dispersion results…

Benchmarks for LOFAR data (SKA Pathfinder)

Results (Fermi Generation)…

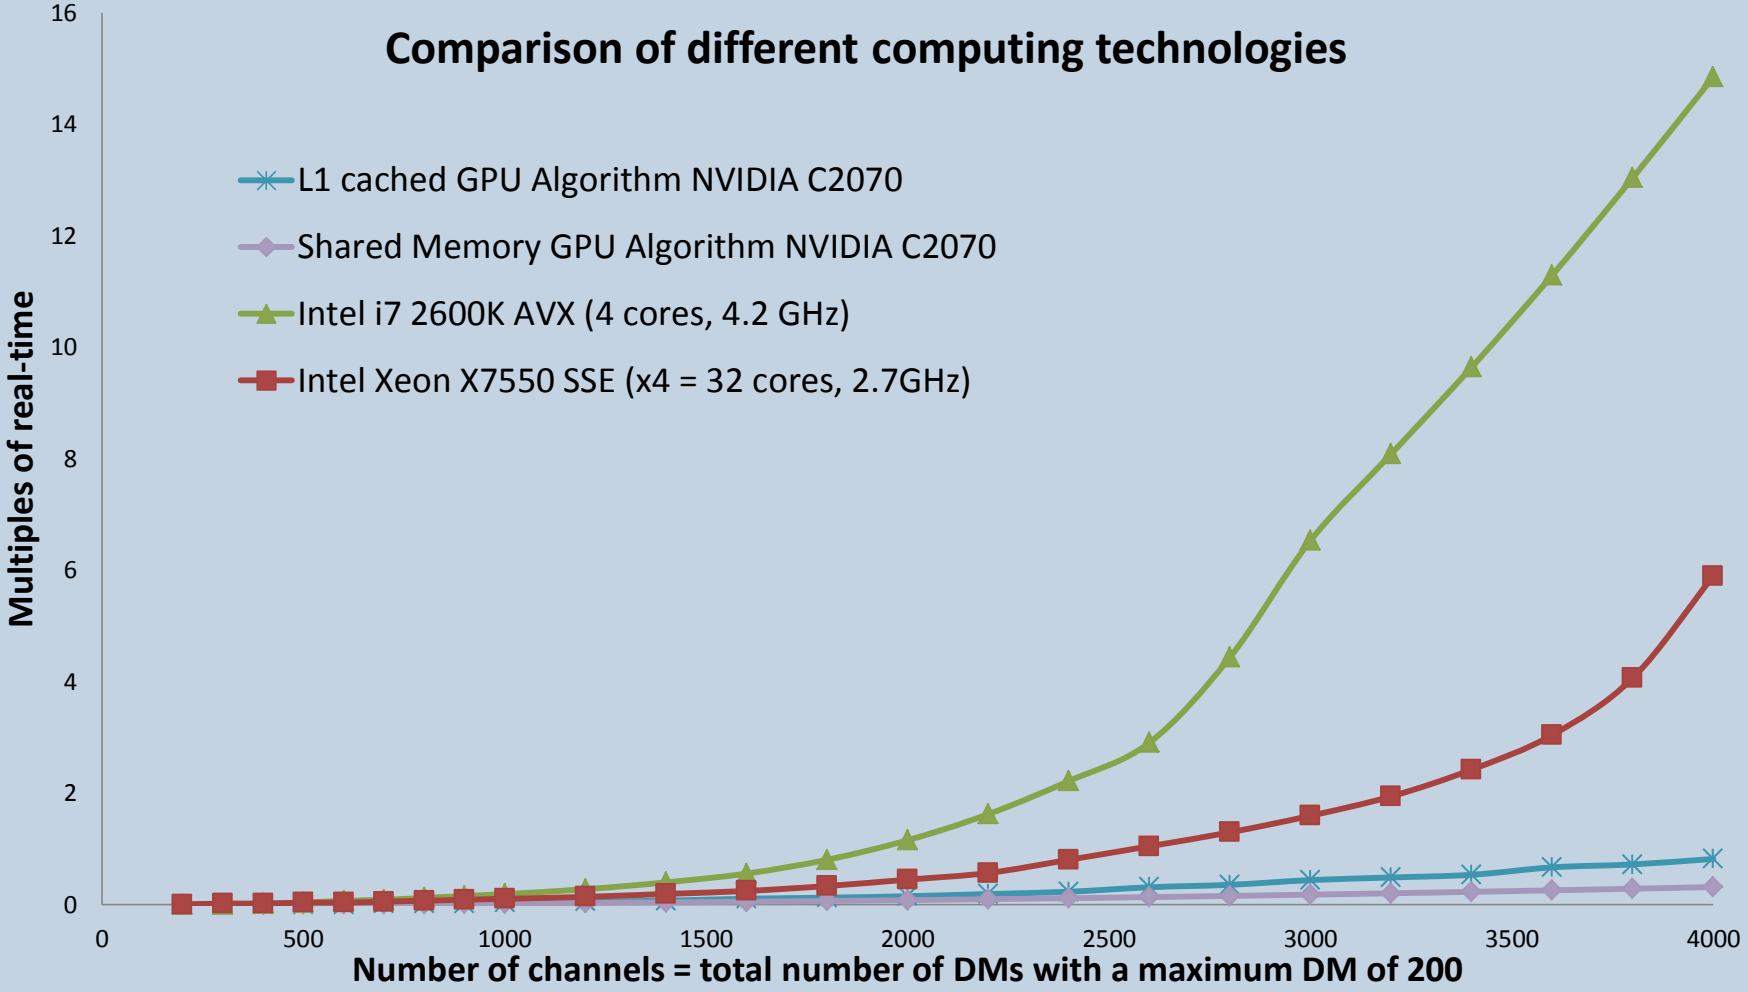

OXFORD e-Research CENTRE

Results (adding more CPU cores)…

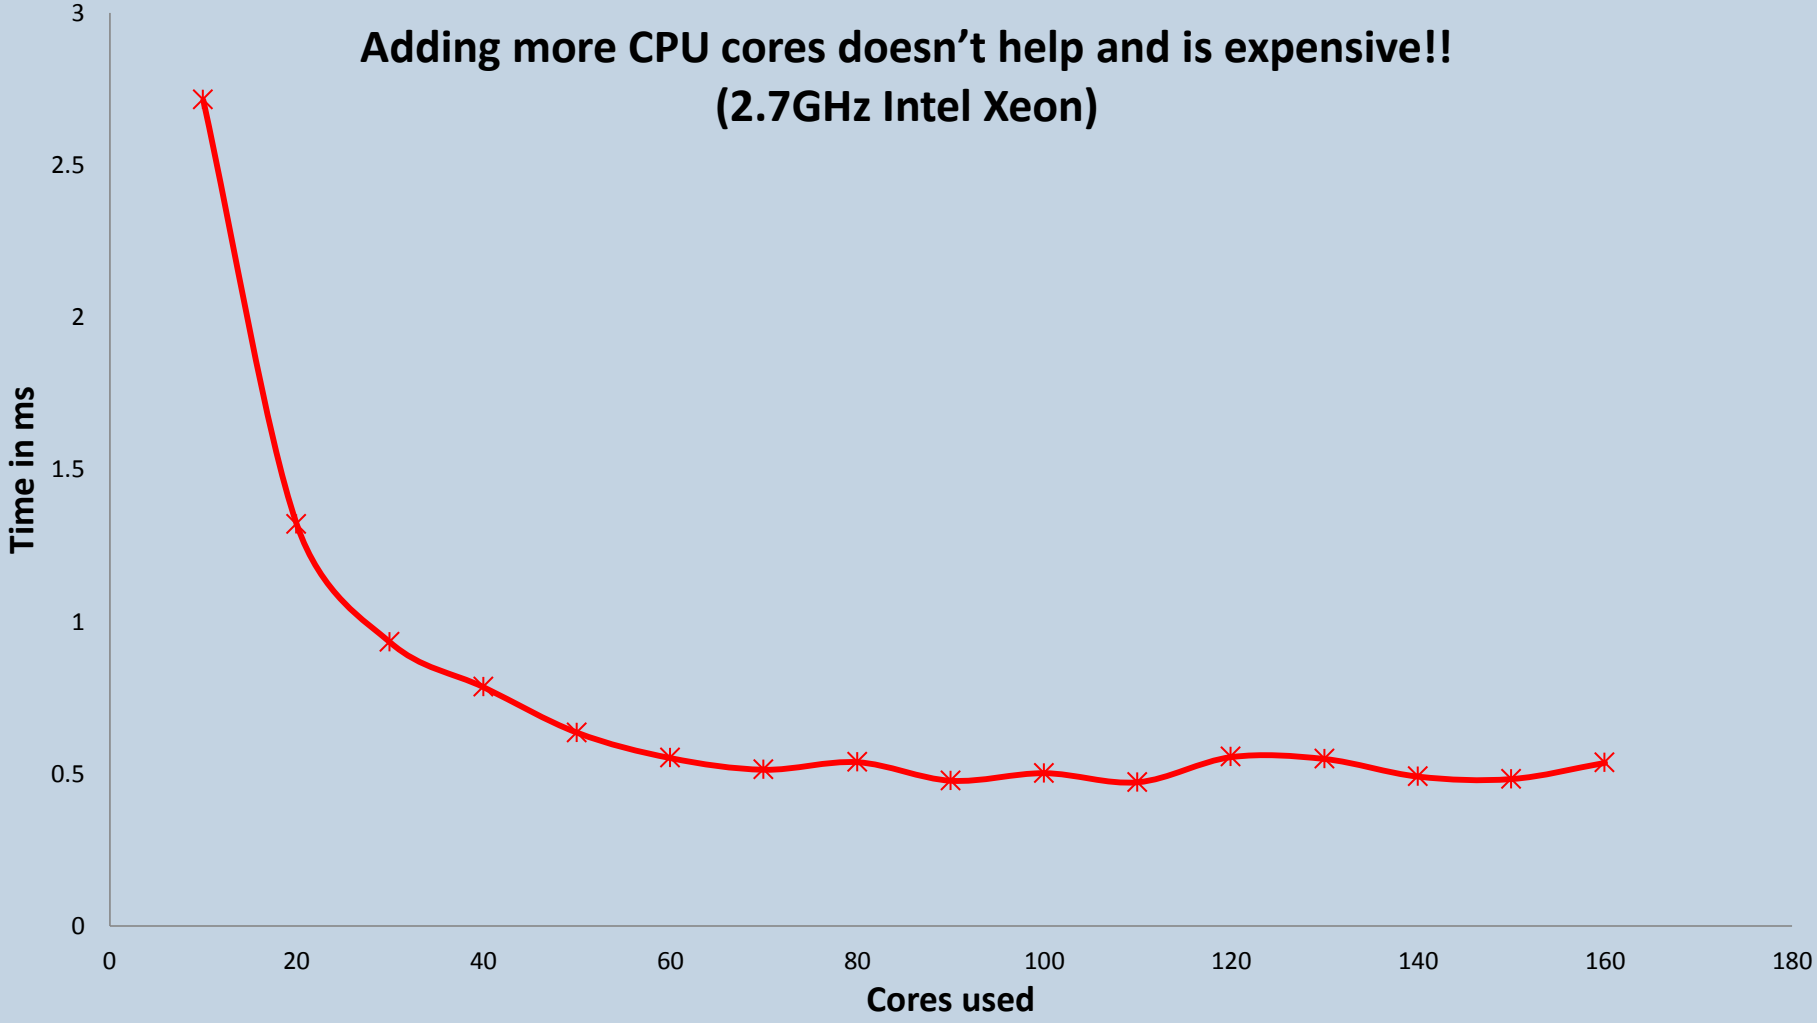

OXFORD e-Research CENTRE

Results (Fermi and Kepler)…

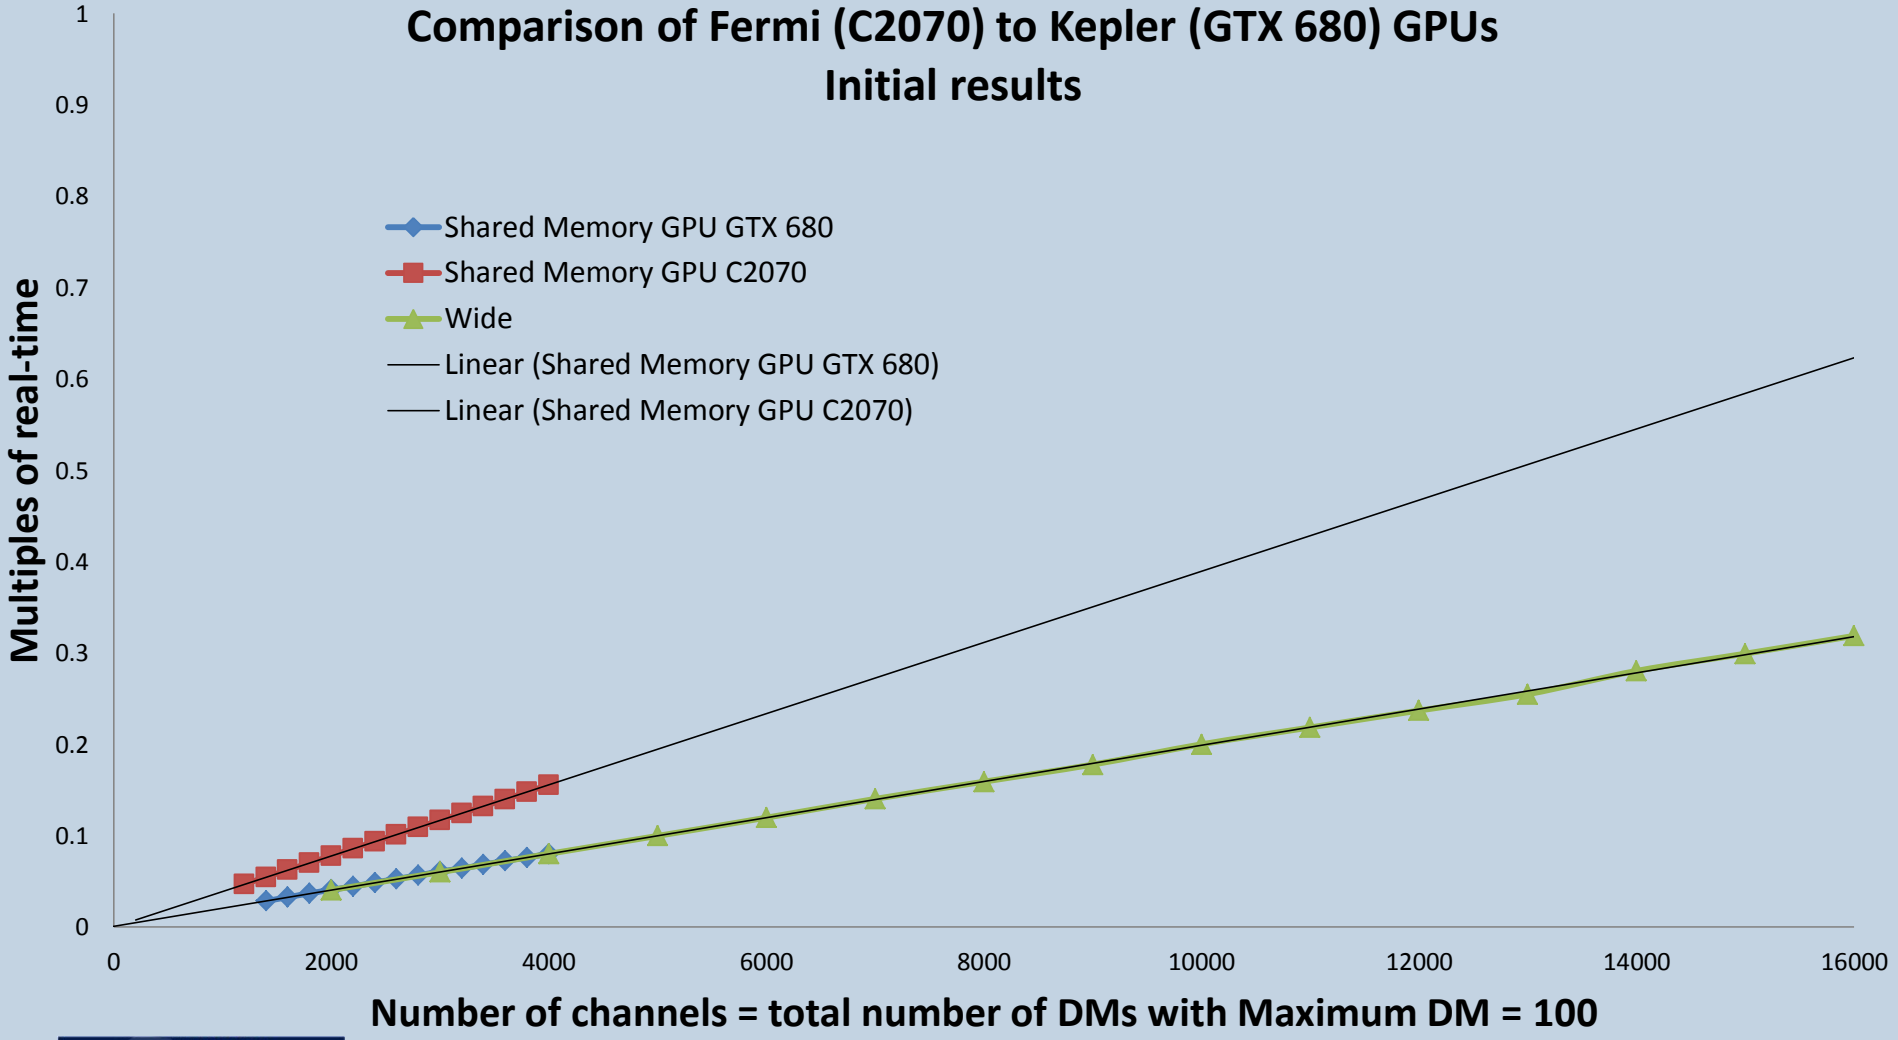

OXFORD e-Research CENTRE

Results (Kepler data paths)…

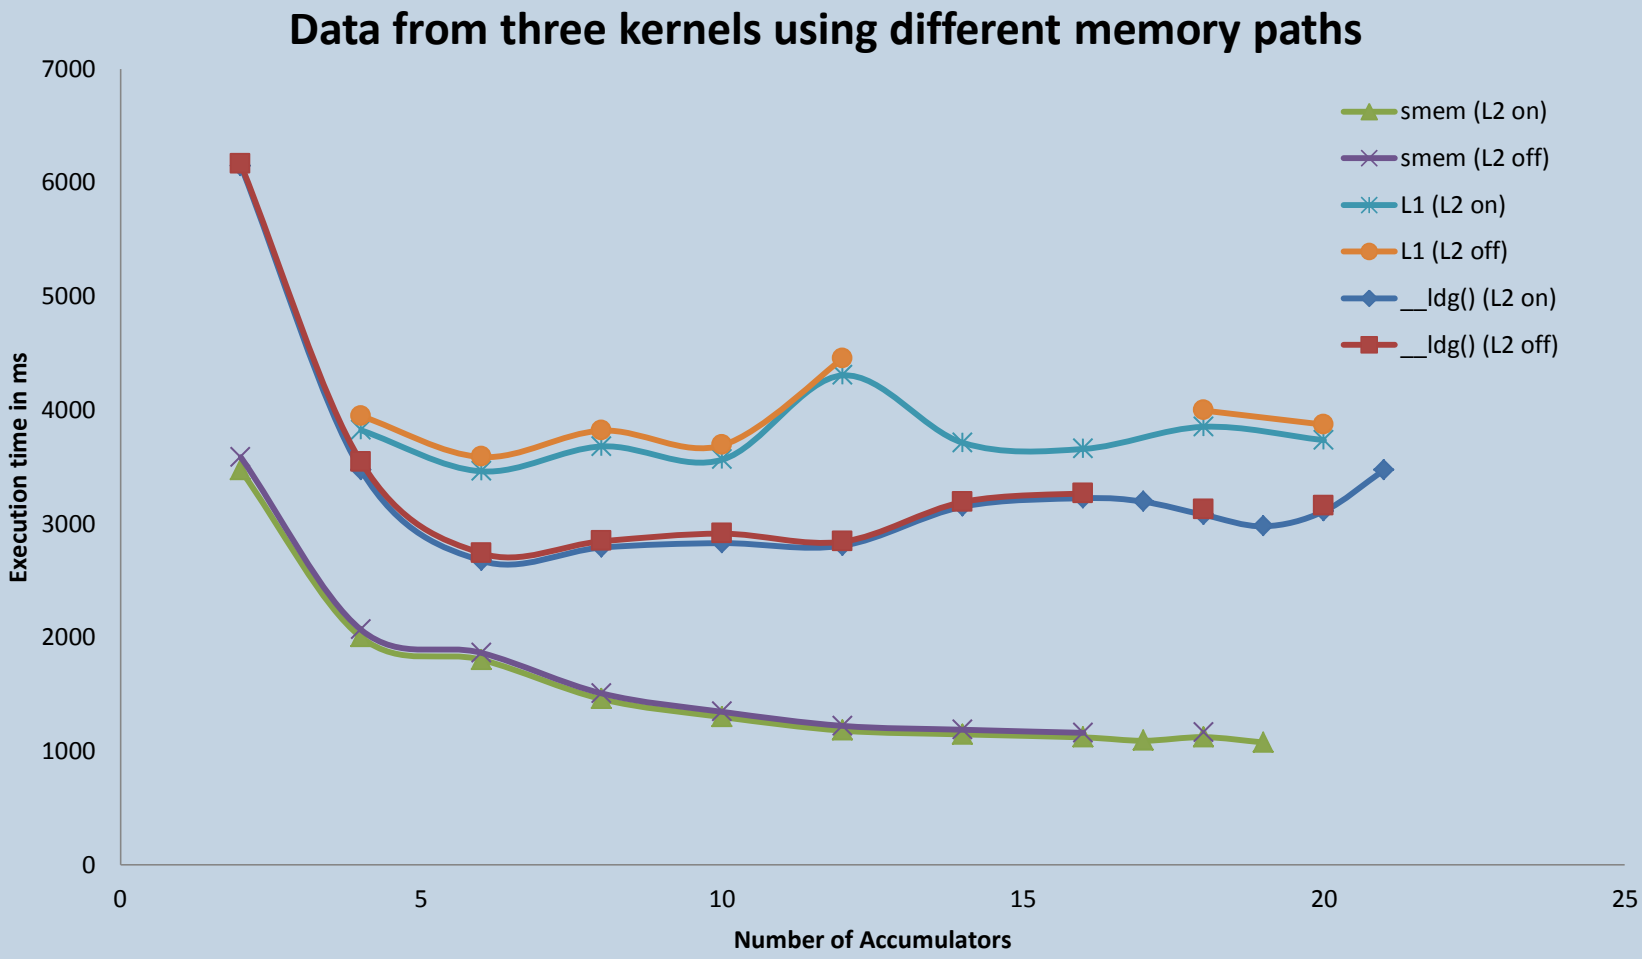

De-dispersion results…

Application to SKA?

De-dispersion results…

4000 Frequency channels
50 microsecond time sampling
300MHz bandwidth
1400MHz Central frequency
8 bit data
6700 trial DM searches

DM Low	DM High	DM Step	# DM's	Down-sample
0.000000	594.000000	0.180000	3300	1
594.000000	1071.000000	0.360000	1325	2
1071.000000	2097.000000	0.720000	1425	4
2097.000000	3033.000000	1.440000	650	8

De-dispersion results…

Using my smem code, Binning in (f,t) at 2x diagonal DM (No double buffering)

Code achieves 65-70% of peak shared memory bandwidth

GPU	Kepler (K20c - current)	Volta (Gemini configuration - predicted)	Volta with above optimisations (predicted)
Fraction of real-time	1.37	0.2 – 0.4	0.15 - 0.3
Number of GPUs needed (per beam)	2	0.2 – 0.5	0.15 - 0.333
Watts per beam (Max)	450 W	45 - 90 W	30 - 80 W
Cost per beam (capital, GPU only)	£6 – 8K	£700 - 1500	£ 500 – 800
Cost per beam (running, 2 year survey, GPU only)	~£1K – 1.5K (based on 1KWh costing £0.2)	~£150 - 300 (based on 1KWh costing £0.2)	~£100 - 200 (based on 1KWh costing £0.2)

De-dispersion results: Hardware Tuning…

Configuration	GPU , PCIe, disk access and analysis	GPU and PCIe only
GTX680 + SATA (optimised)	2.53	1.92
GTX 670 + SSD (un-optimised)	3.7	3.33
GTX 670 + SSD (optimised)	2.53	1.92

- Custom kernel
  - CentOS/RHL 3.10.18
- SSD disks
  - I/O Tuning
- NVIDA 331.17 Driver
  - Nvreg_EnablePCIeGen3=1

With Rahim (Raz) Lakhoo rahim.lakhoo@oerc.ox.ac.uk (U. Oxford)

On-going and Future Work…

- Execution of standalone code is dominated by analysis on CPU – Investigate moving this (or part of) to GPU.
- Work on a GPU polyphase filter module is underway (Jan Novotny & Karel Adamek– Opava).
- Some preliminary work aimed at producing a GPU version of the ARTEMIS RFI-Clipper has been undertaken (Patrick Hollebon – NVIDIA funded Student).
- 2x PDRA to begin in March 2014 to further the work, Long 1D FFTs, Harmonic Sum…
- Paper almost submitted, code will be on git soon!

Acknowledgments and Collaborators

Astro-Accelerate : http://www.oerc.ox.ac.uk/research/wes

ARTEMIS : http://www.oerc.ox.ac.uk/research/artemis

University of Oxford

Mike Giles (Maths) – Cuda, GPU algorithms.

Aris Karastergiou (Physics) – ARTEMIS, Astrophysics, Experimental Work.

Chris Williams (OeRC) – RFI Clipper, Data pipeline.

Steve Roberts (Engineering) – Signal processing/detection algorithms.

University of Bristol

Dan Curran (Electrical Engineering) – OpenCL work.

Simon McIntosh Smith (Electrical Engineering) – OpenCL work.

University of Manchester

Ben Stappers - Testing

Even Keane - Testing and Data

Cees Bassa - Testing and Data