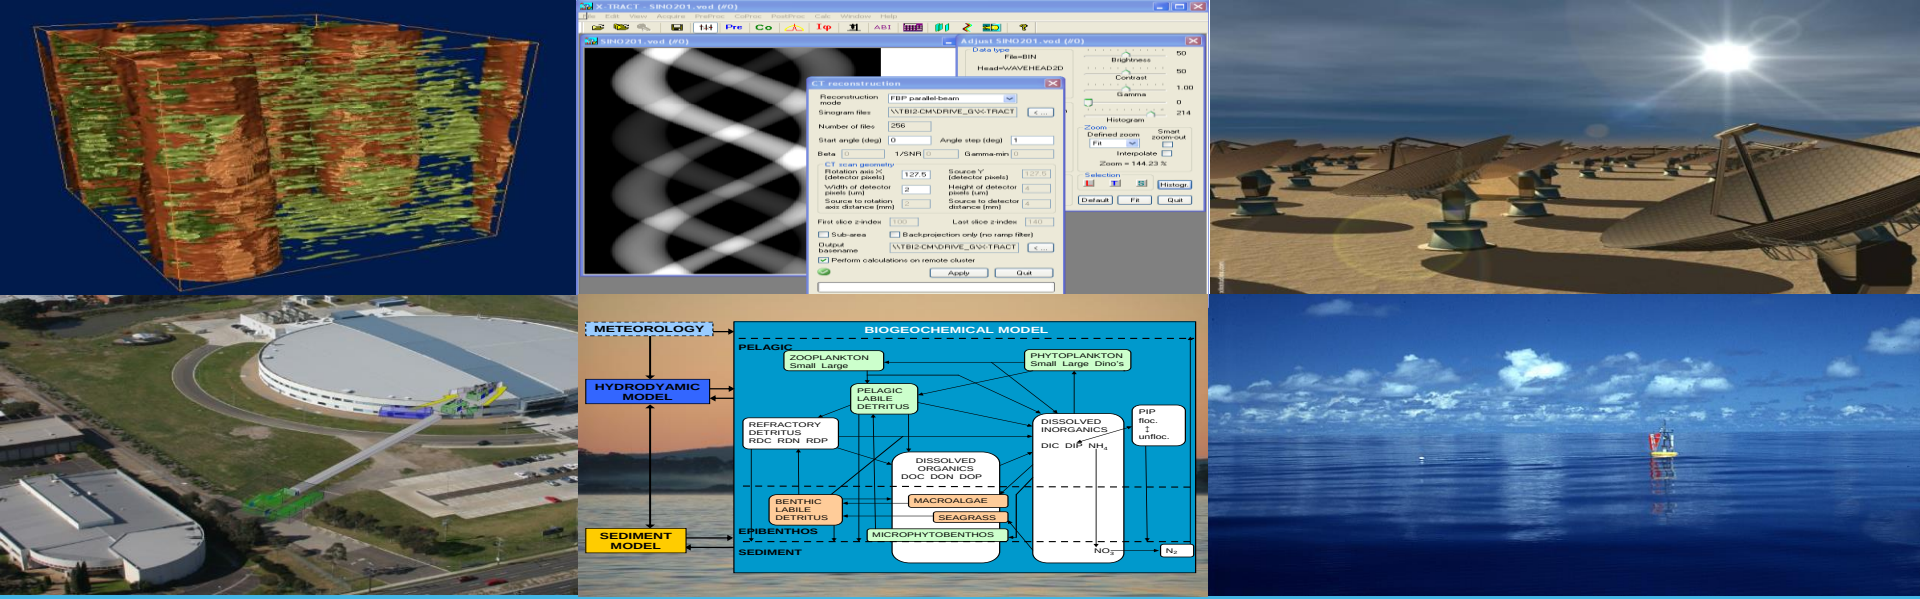

CSIRO Computational Science

Breakthrough Science on GPUs

John Taylor | Director, CSIRO eResearch & Computational and Simulation Sciences

July 2014

About CSIRO

People	5500
Locations	58
Flagships	9
Budget	$1B+

Top 1% of global research institutions in 14 of 22 research fields

Top 0.1% in 4 research fields

62% of our people hold university degrees

2000 doctorates

500 masters

In partnership with universities, we develop **650** postgraduate research students

2009: CSIRO Bragg Cluster Launch, first of its kind in AU

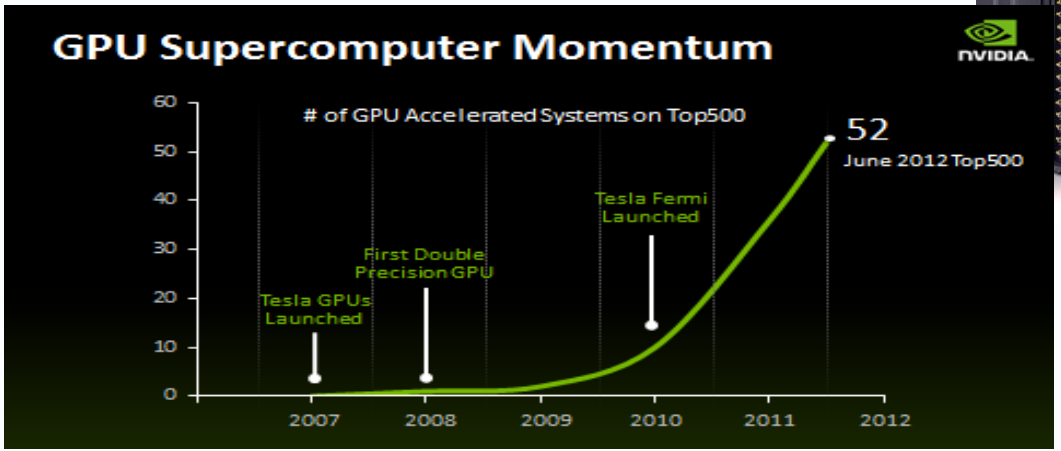

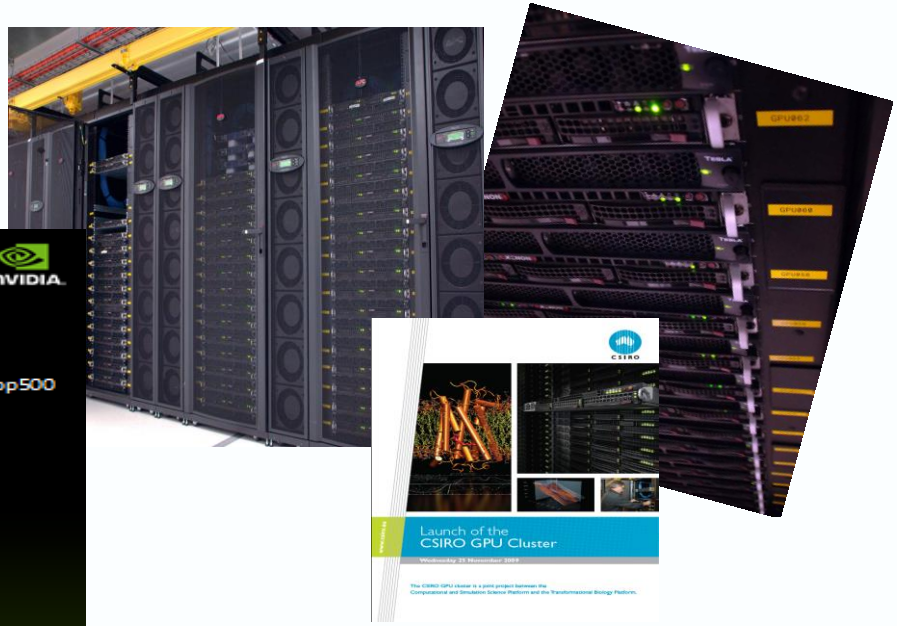

2013: Bragg upgrade to latest generation GPU, 384 Kepler K20M GPUs

June 2014:
#181 TOP500 List
#8 Green500 List

CSIRO Bragg GPU Cluster

TOP500 and Green500 Rankings

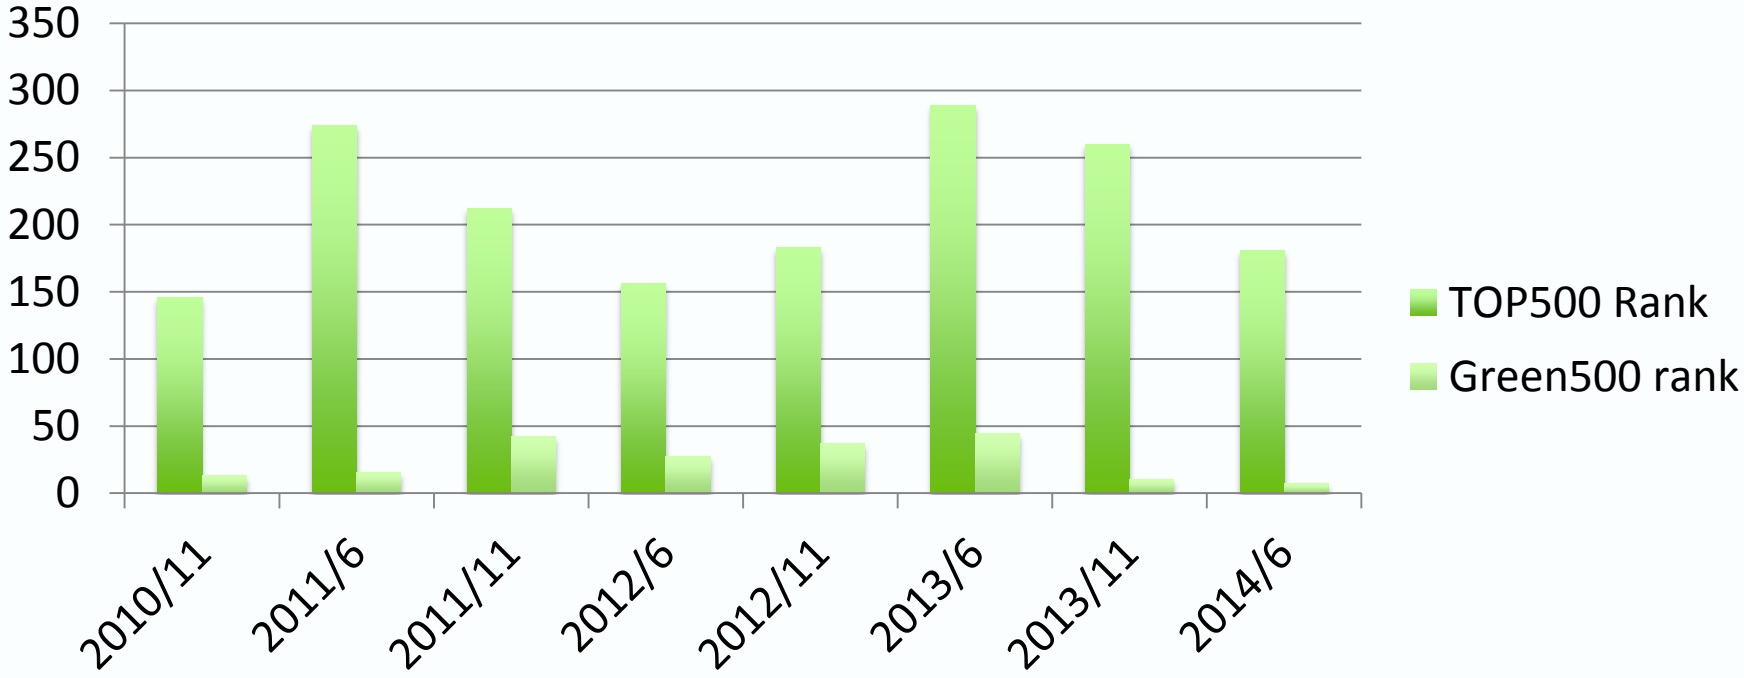

CSIRO Bragg GPU Cluster

TOP500 Performance

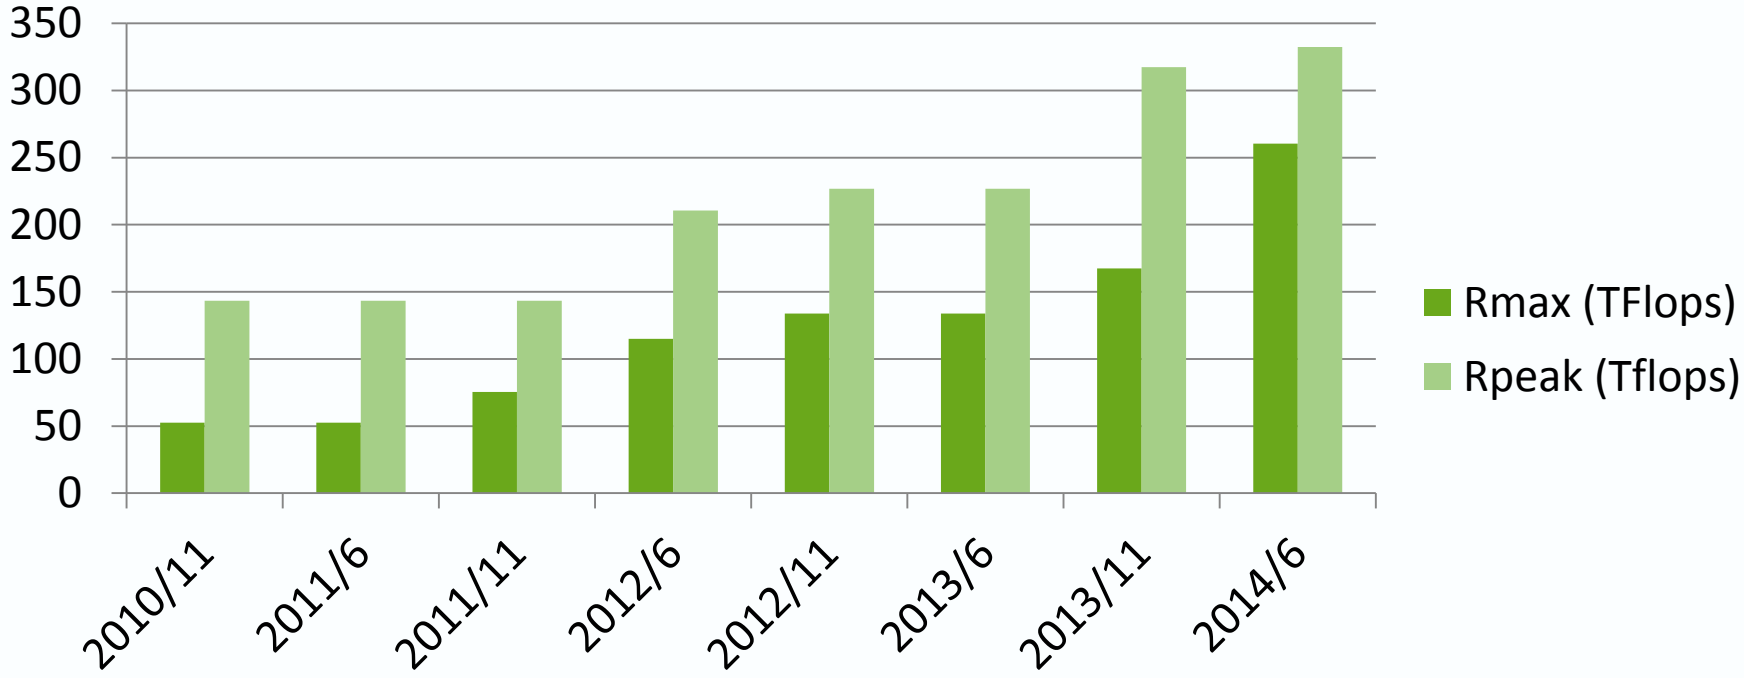

CSIRO Bragg GPU Cluster

Peak performance from CPU and GPU

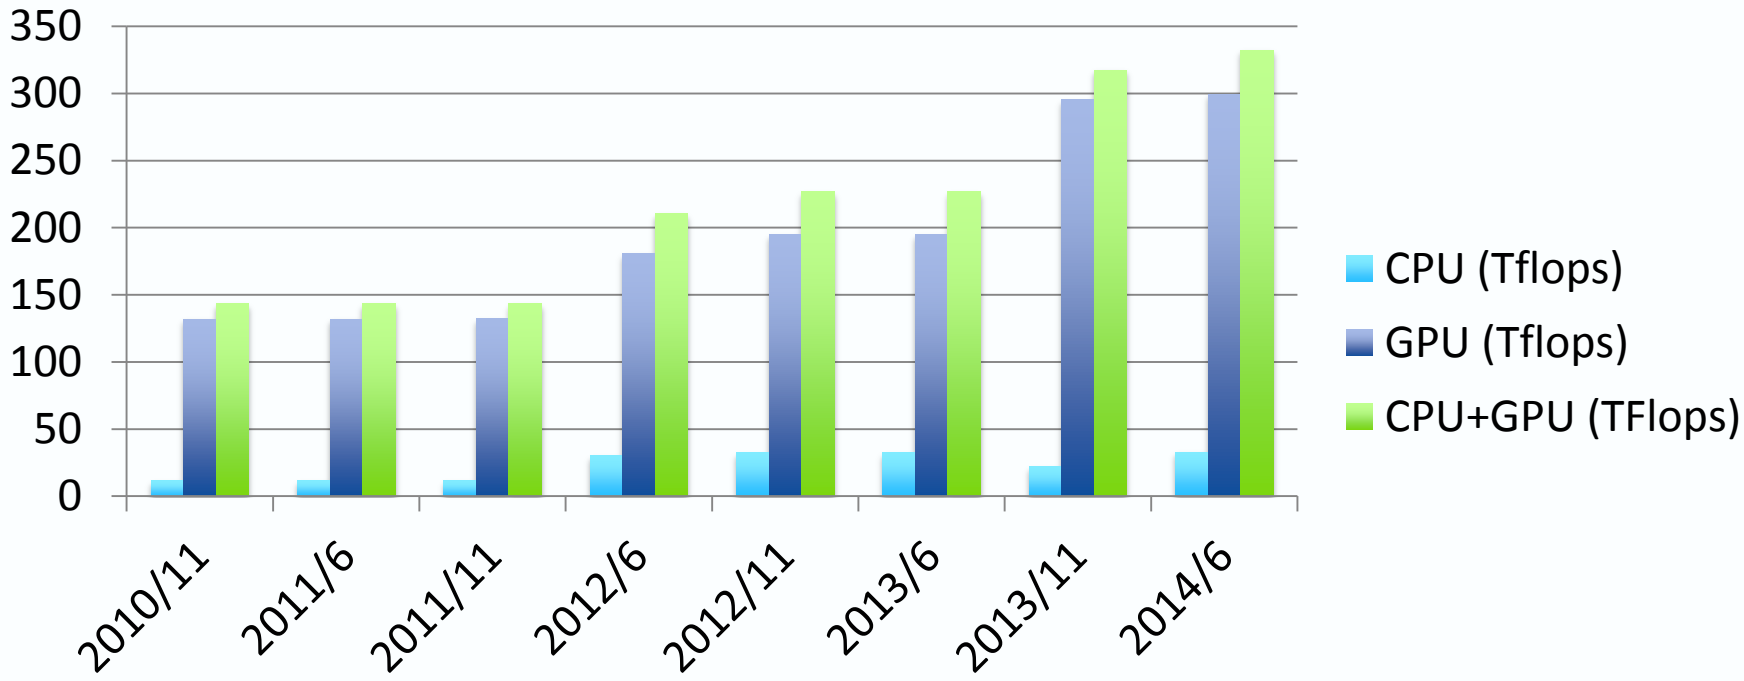

CSS Capability Development Platforms

CSIRO Computational & Simulation Sciences

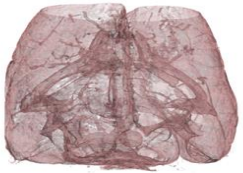

Imaging Processing & Visualisation

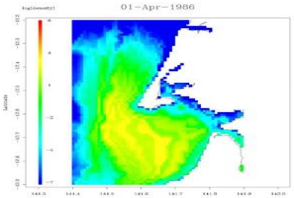

Model Data Fusion & Spatial Modelling

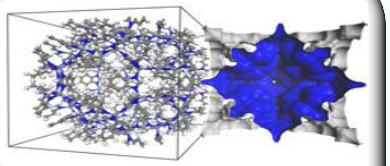

Computational Material Design

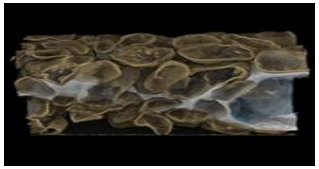

Data Constrained Materials Modelling

- Developing core, domain specific, computational capability
- Developing generic computational tools to aid domain work across CSIRO
- Deploying GPU capability/tools back into Flagships

Principles of X-ray Computed Tomography (CT)

- Typical experimental set-up involves an X-ray source, a sample on a rotation stage and a 2D position-sensitive detector
- Images of the sample are collected at many different rotation angles spanning 180 or 360 degrees
- Acquired images are processed in a computer to produce a 3D representation of the internal structure of the sample

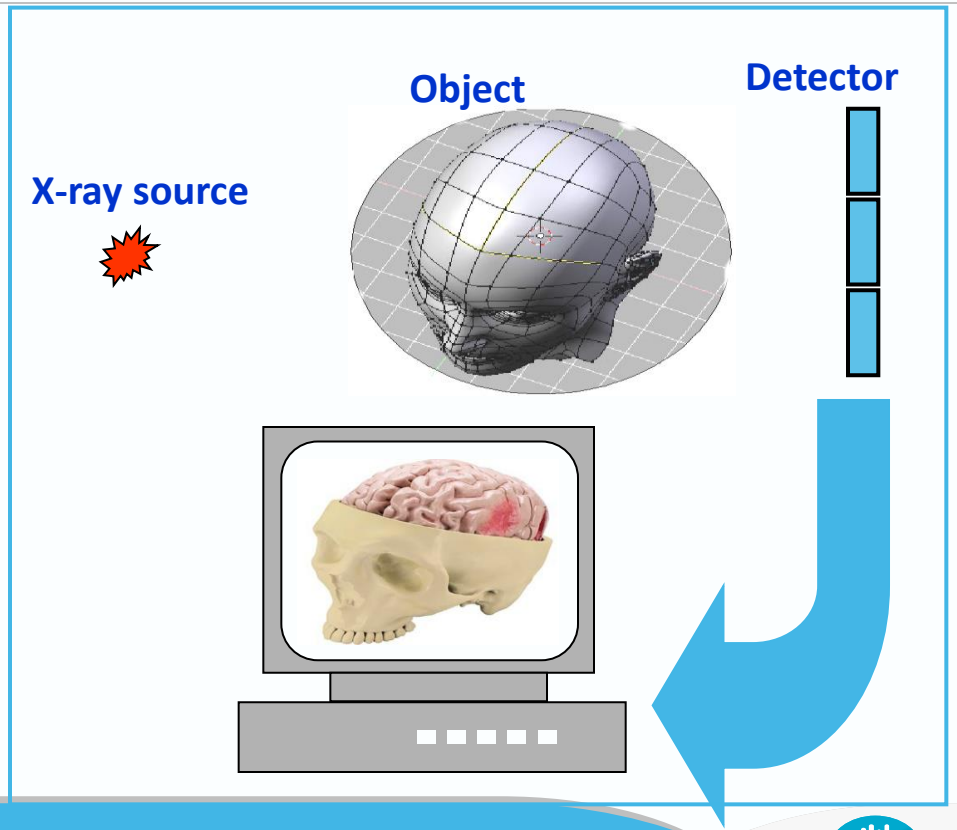

X-TRACT synchrotron map

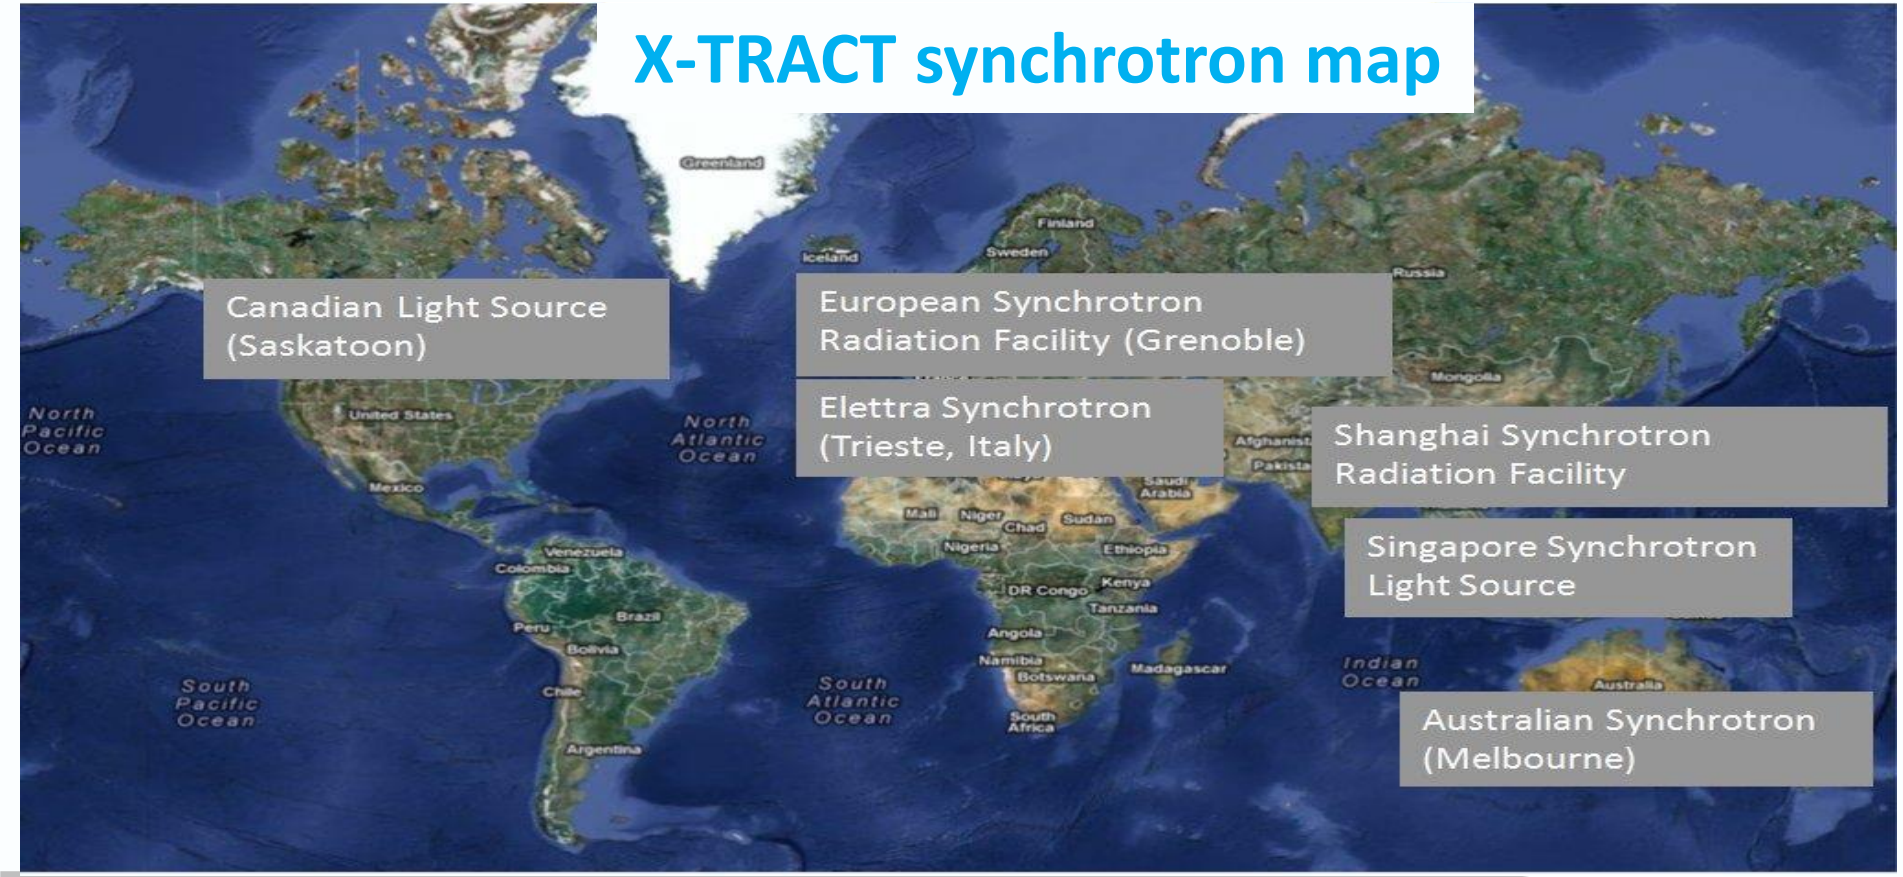

Radiation therapy applications

Modern radiation therapy is to a large extent a computational discipline and can greatly benefit from use of task- and data-parallelism. Some applications were demonstrated on GPUs already:

- CT reconstructions
- Image registrations
- Treatment planning
- Dose computations (e.g. X Gu, U Jelen et al 2011 PMB 56)

Need for speed: imaging and treatment verification can be used as feedback to improve the treatment (adaptive radiotherapy), currently offline (mostly population-based), one day online.

Particle (proton/carbon ion) therapy with raster scanning @ University of Marburg:

- most precise external beam technique (only 5 centers worldwide: 3 active, 2 to start)
- increased precision = increased need for verification (more computations)
- longer computational times (small head case: 1 hour on single-thread)

Collaborative project between CSIRO and University of Marburg

Plan robustness in radiation therapy

Automatic discovery of robust beam setups.

Results (mean and sd for a single beam):

- 4-core Intel Xeon W3530 2.8GHz 12GB RAM + NVIDIA Tesla C2050 3GB RAM
- 10 skull base cases, 42 beams directions (10 runs each for timing stats)
- 4k-40k pencils of 120-350 samples, 2 mm analysis radius (0.5 mm step)
- Single-precision floating-point operations only (sufficient precision)

mean(sd) ms	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	Pool
Native (1 thread)	21299 (6628)	9891 (2837)	6258 (1485)	15768 (4959)	4342 (1136)	10888 (3179)	10117 (2849)	5464 (1470)	8155 (2195)	11388 (3936)	10357 (5941)
GPU OpenCL	219 (109)	122 (51)	88 (38)	148 (56)	61 (24)	160 (65)	151 (64)	52 (22)	109 (46)	126 (61)	124 (75)
Gain	**119 x** (36)	**98 x** (34)	**87 x** (30)	**123 x** (36)	**83 x** (25)	**81 x** (24)	**82 x** (30)	**124 x** (42)	**90 x** (31)	**106 x** (29)	**99 x** (36)
CPU OpenCL	6498 (1996)	2552 (615)	1898 (438)	4810 (1495)	1324 (331)	3280 (944)	3051 (841)	1396 (310)	2481 (649)	2935 (818)	3022 (1798)
Gain	3.3 x (0.0)	3.8 x (0.4)	3.3 x (0.0)	3.3 x (0.0)	3.3 x (0.0)	3.3 x (0.0)	3.3 x (0.0)	3.9 x (0.4)	3.3 x (0.0)	3.8 x (0.4)	3.5 x (0.3)

Cloud Based Image Analysis and Processing Toolbox

Available now → www.cloudimaging.net.au

CloudBased
Image Analysis & Processing Toolbox

Image Analysis and Processing for *for everyone.*

The Cloud-based Image Analysis and Processing Toolbox project provides access to existing biomedical image processing and analysis tools via remote user-interface using the NeCTAR cloud.

Use Toolbox | Demo | Project Blog

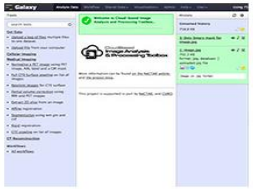

Use project's free server

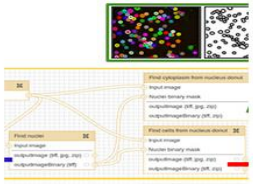

Watch other demos

Project Blog

Create and Share your imaging workflows with other scientists.

Connect various imaging functions to create more complex imaging algorithms.

Visualise the results.

Reuse the workflows on different datasets.

Adjust the parameters whenever required.

Open for Australian scientists.

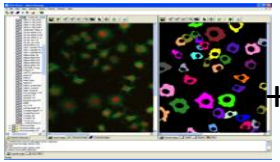

HCA-Vision

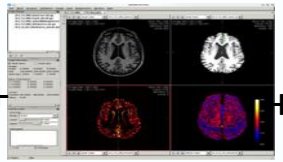

MILXView

X-TRACT

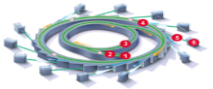

Workflows in action

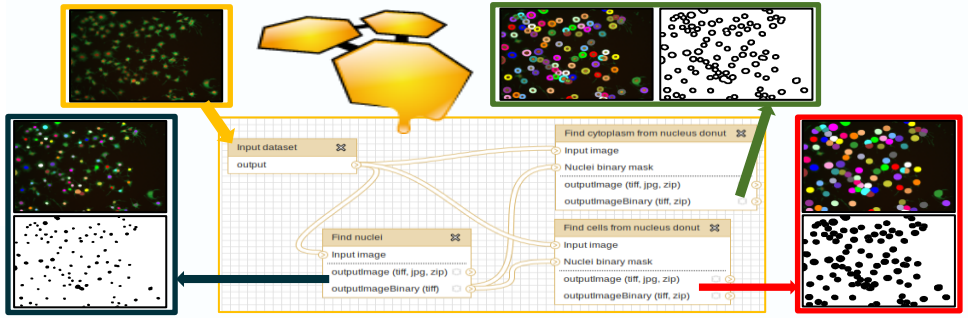

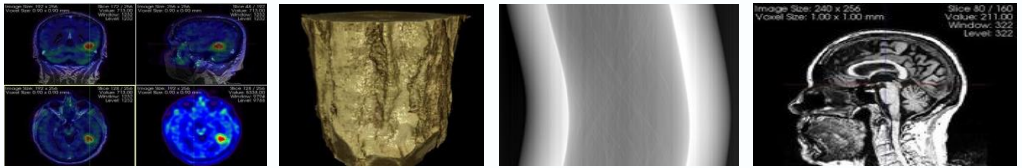

Quantifying Uncertainty and in Soil Carbon Dynamics

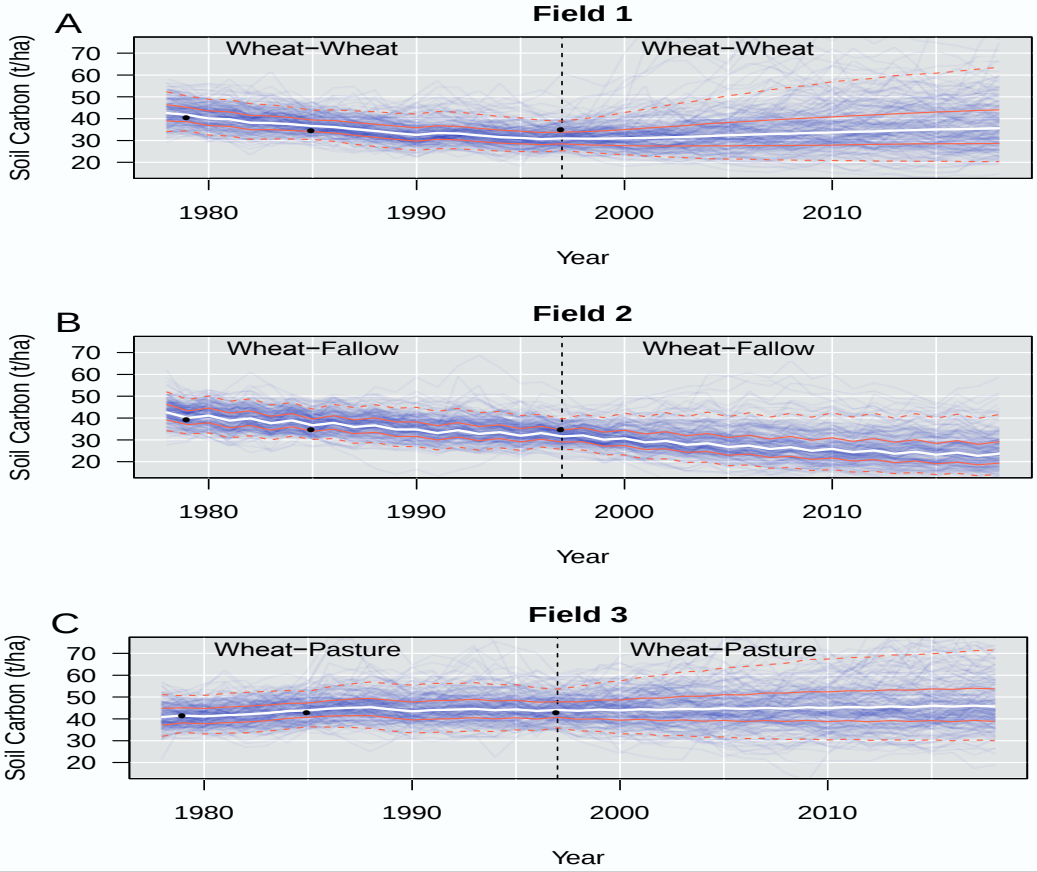

Quantifying sediment Loads to the Great Barrier Reef

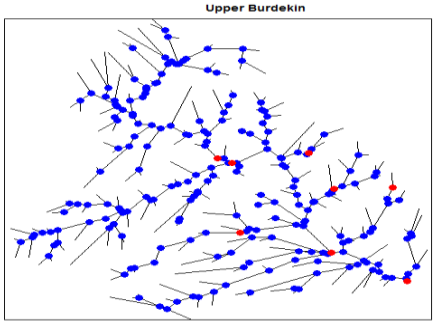

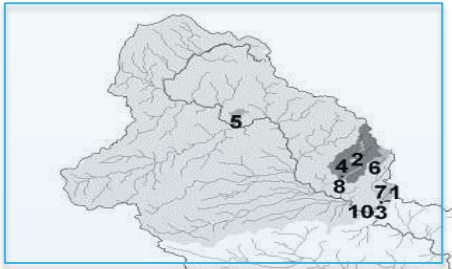

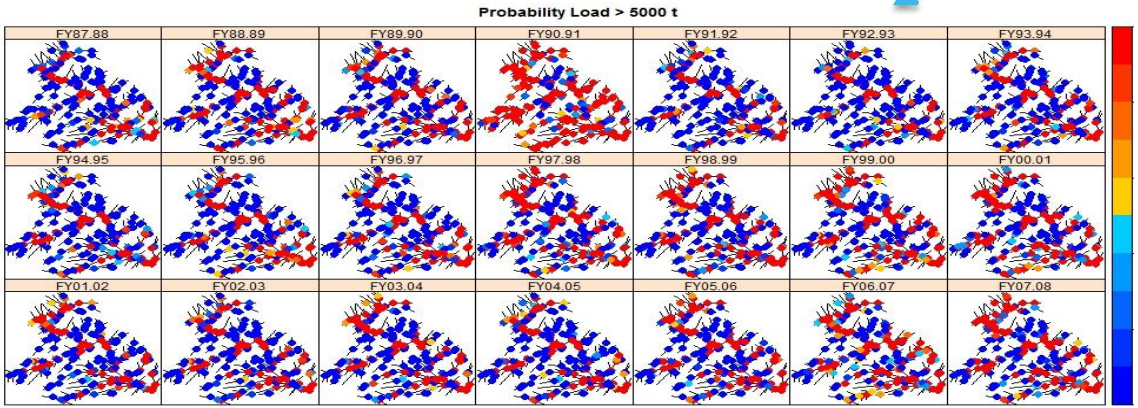

Simulated Nanostructure Assembly (SNAP)

Adding Distributed Parallelism to Support Larger Simulations

Nanoparticle assembly modelling

- Mesoscale particle descriptions (~100 - 1000 nm)
- Original serial code had performance limitations

One project developed a GPU port of SNAP

- 3 to 5x speedup over original code
- Limited GPU memory was restricting the number of particles in a simulation

A follow-up project addressed the simulation size limitations

- Optimisation of GPU memory layout
  - Enables simulations 25 times larger on each GPU
- MPI used to distribute the simulation over multiple compute nodes for tackling even larger problems

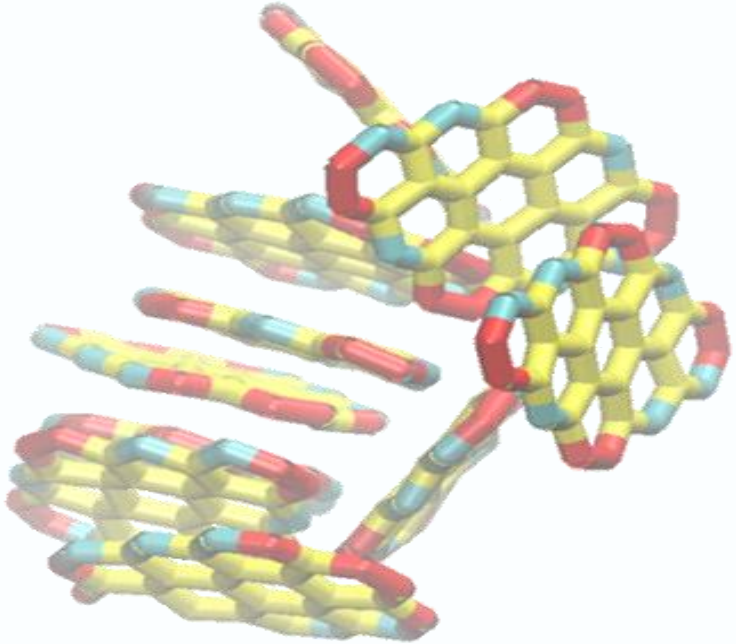

GPU Accelerated Meshing Workflow

CSIRO's Workspace Framework coordinates distributed execution of OpenCL-based meshing operations

www.csiro.au/workspace

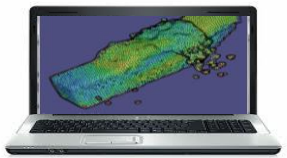

Researcher

VNC with VirtualGL

Workspace UI
IM&T Remove Viz

CSIRO GPU Cluster

Workspace Distributed Execution
(Via PBS Job System)

OpenCL GPU Accelerated Workspace Operations

Surface Mesh Generation for Particle Simulations

Improving Performance with OpenCL and GPUs

3D fluid dynamics particle simulations

- Simulations approaching 100 million particles

Changed to per-particle volumetric calculations

New OpenCL atomic operations were effective

- Integer based atomic_max() used with floating point data
- ~10x faster than memory area locking code
- ~5x speedup overall

Future Work - Should be scalable to multiple GPUs

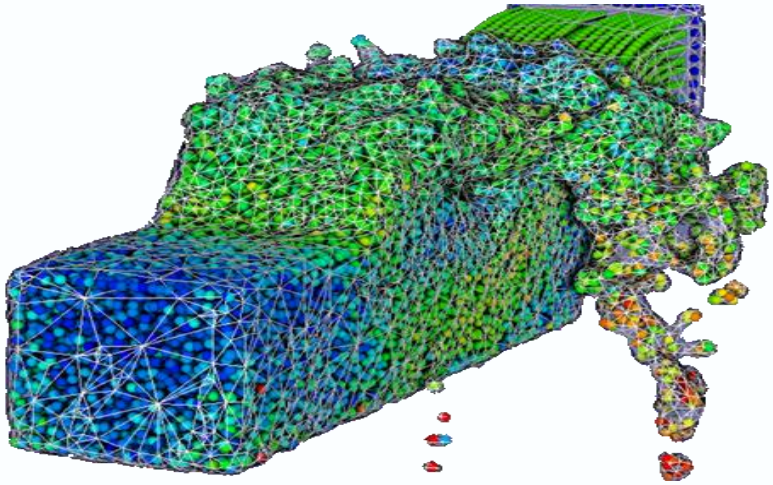

AWRA Requirements:

- C# codebase
- AWRAkernel – independent computations on a per pixel basis, with lots of computation
  - Indicated GPUs should be useful for calculations
- Required max() min() and exp()/log() transcendental functions to be available
- Able to be run on CPU for Bureau of Meteorology implementation
- Possibility of integration with R research code base.
- OpenCL meets these requirements

Timing - Data movements and Kernel computation

- AWRA DA kernels computation and data transfer
- Three kernel computations required per time step taking approx. 45% of original time.
- Murrumbidgee catchment dataset, single precision.
- Original C# single core implementation of “Kernel Computation” takes around 1.5 seconds using W5590 CPU.

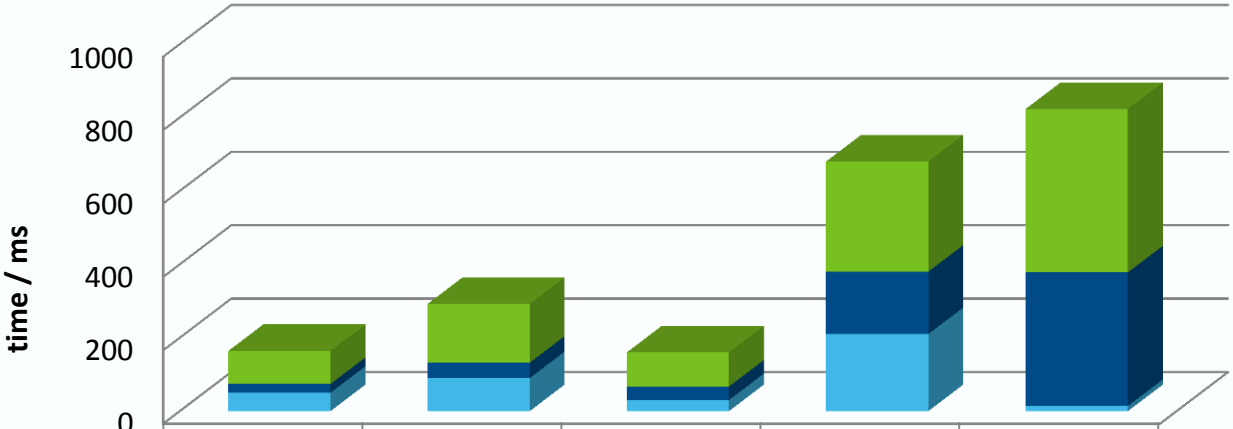

	AMD 6950 GPU / 3.3GHz CPU	NVidia C1070 GPU / 3.3 GHz CPU	NVidia K20m GPU / 2GHz CPU	Intel HD4400 GPU / 1.9GHz CPU	CPU W5590 4 cores 3.3GHz (OpenCL)
Inflate, transfer and compute kernel	90	160	93	300	445
Kernel Computation	24	42	37	170	364
Upload Model Data To Device	50	90	30	210	14

Workflow Services

Infrastructure / Software

- Increase uptake of Workspace and Galaxy
- Utilise CSIRO's supercomputers and Cloud for better turnaround – increased productivity

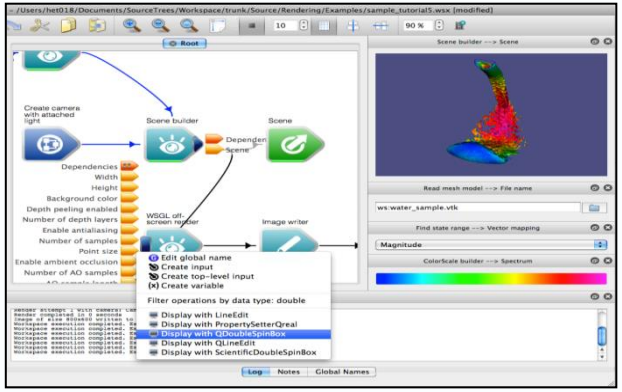

Outreach

- Training and documentation
- Promote and develop culture around workflow management and data provenance

Project Work

- Use existing eResearch Project biannual RFP
- Assist researchers to build high quality workflows

Leverage existing capabilities

- Compute, data, visualisation, CCI, Bioinformatics Core
- Partners

CSIRO Collaboration Platform

The CCP has been successfully licensed to Corporate Initiatives in August 2013

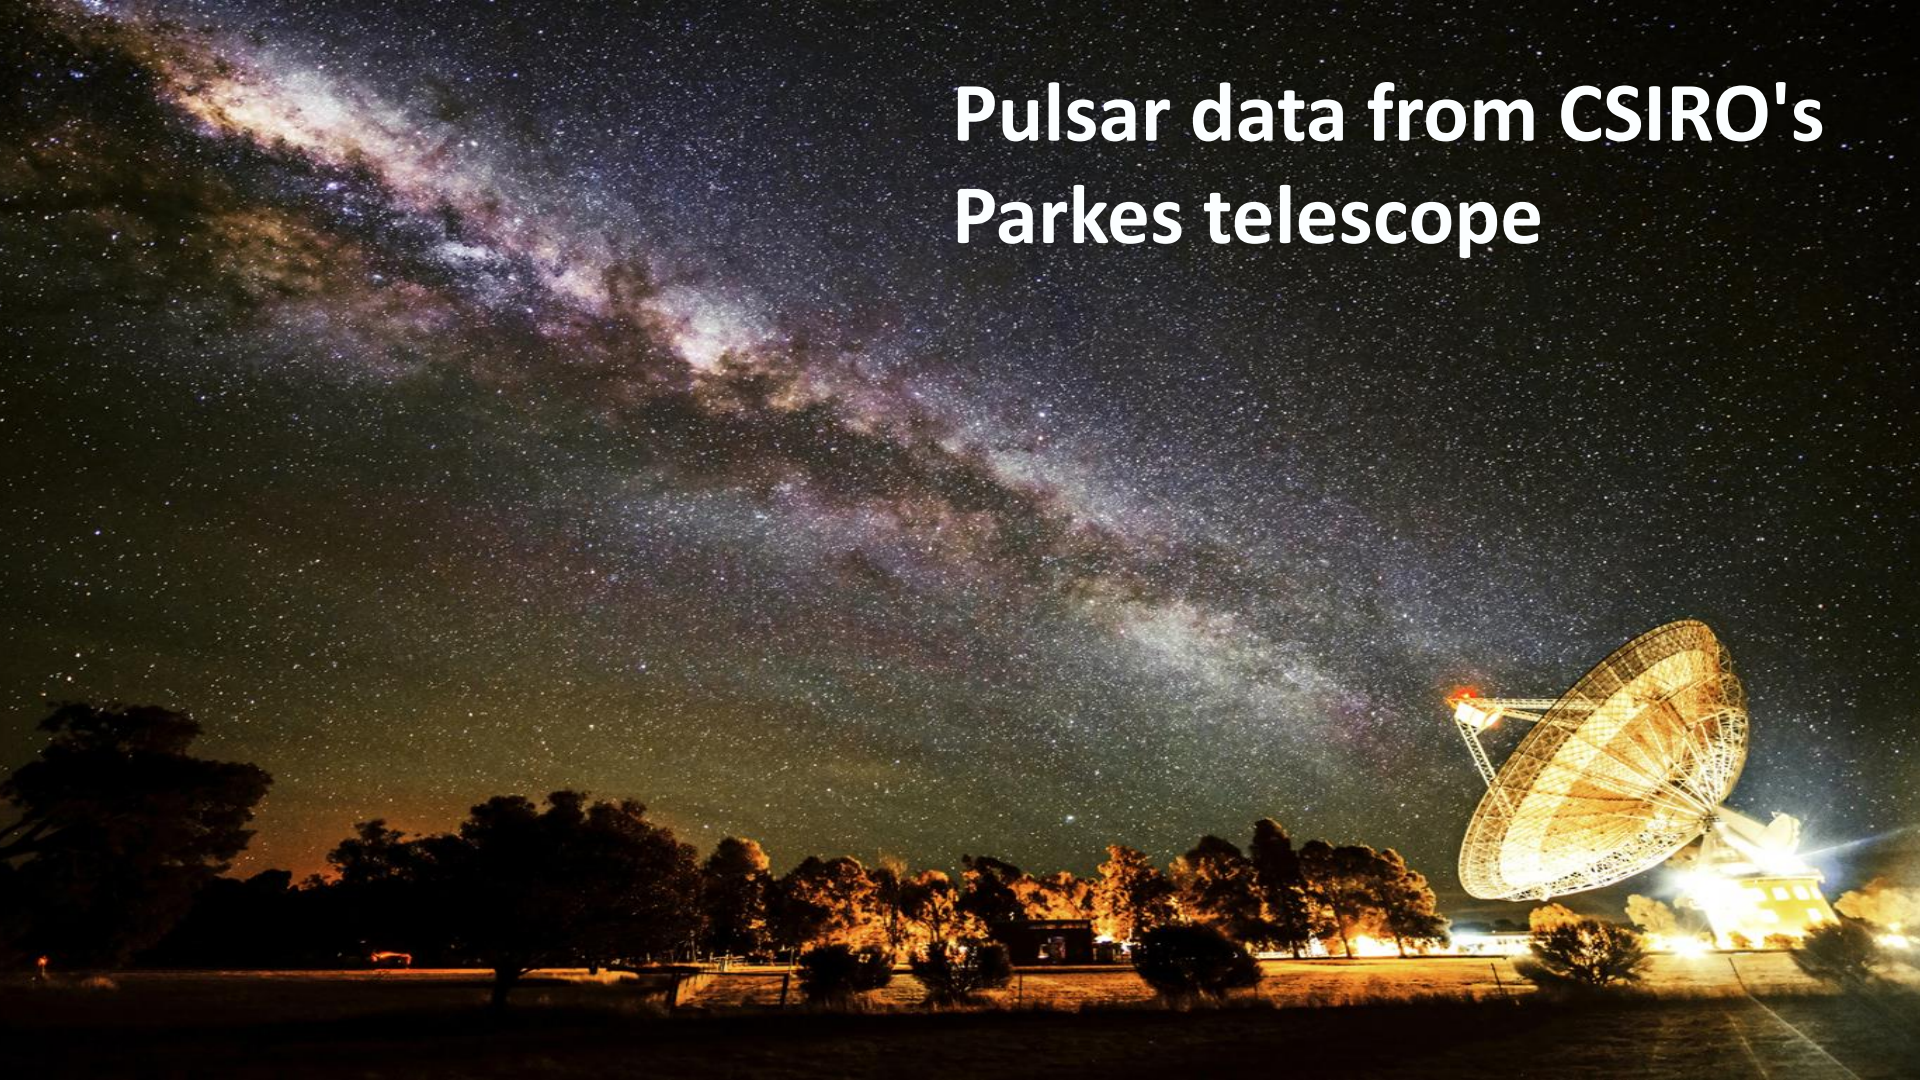
Pulsar data from CSIRO's
Parkes telescope

PPTA-HPC progress to date

Opportunity? Providing external collaborators access to **internationally significant science data** + **compute to process** = "Science as a Service"

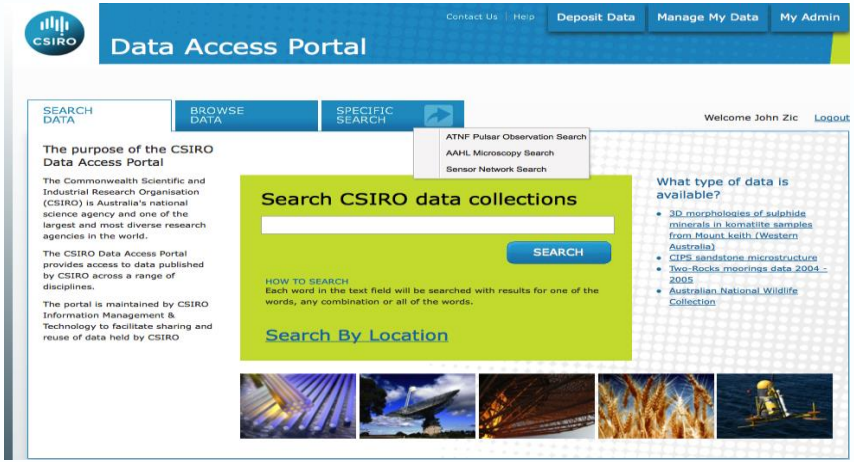

DAP pulsar repository

Compute on Bragg Cluster

Data Source: SilviScan 3 wood properties analysis

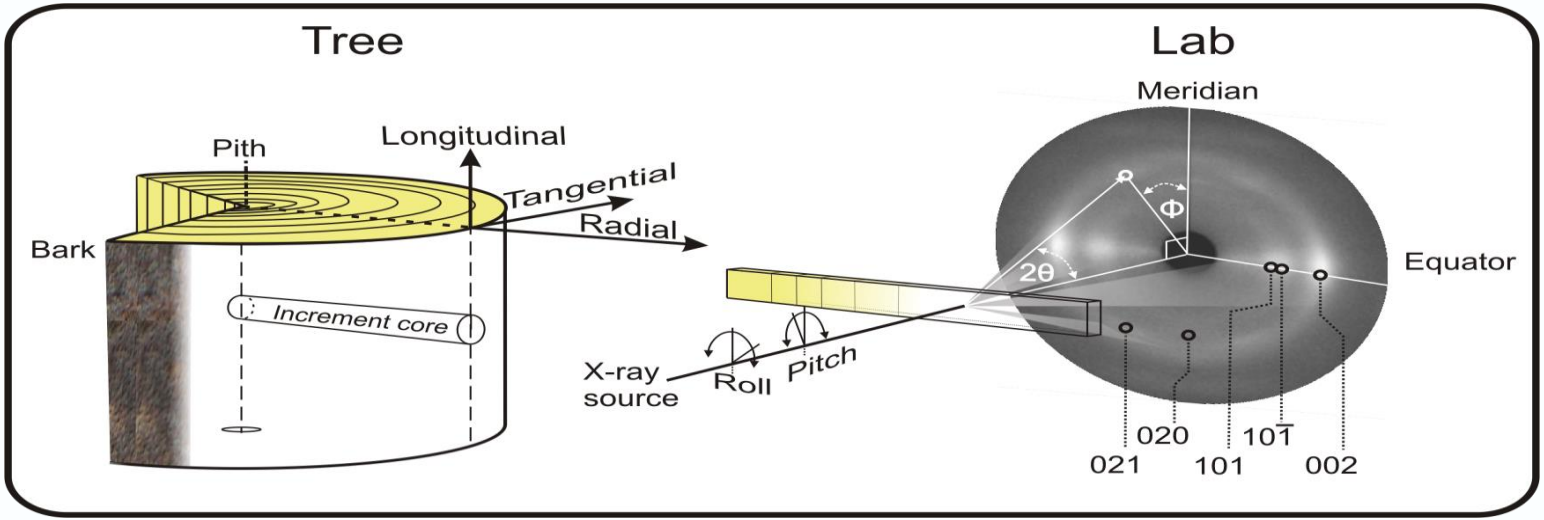

- Diffractometer provides microfibril angle, fiber angular information, crystalinity and cellulose crystal width (work in progress).
- Data requirements
  - High spatial resolution imaging at 200 um steps
  - Can collect ~2000-8000 images per day up to16 GB uncompressed

PCA of larger data sets – Requires multiple distributed GPUs

Statistical methods for interrogating large data sets

- Principal component analysis (PCA):
  - Useful for preliminary and exploratory data analysis
  - Also used in machine learning, outlier detection, regression and prediction procedures
- NIPALS: An iterative method for finding principle components
  - Developed in 1960's by H. Wold

$$\mathbf{M}_{m,n} \xrightarrow{PCA} \mathbf{P}_{m,c} \times \mathbf{T}_{c,n} + \mathbf{R}_{m,n}$$

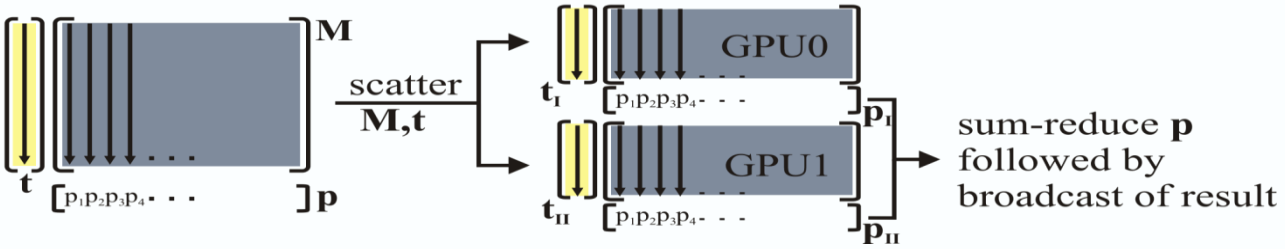

Principal Component Analysis with NIPALS using GPU clusters

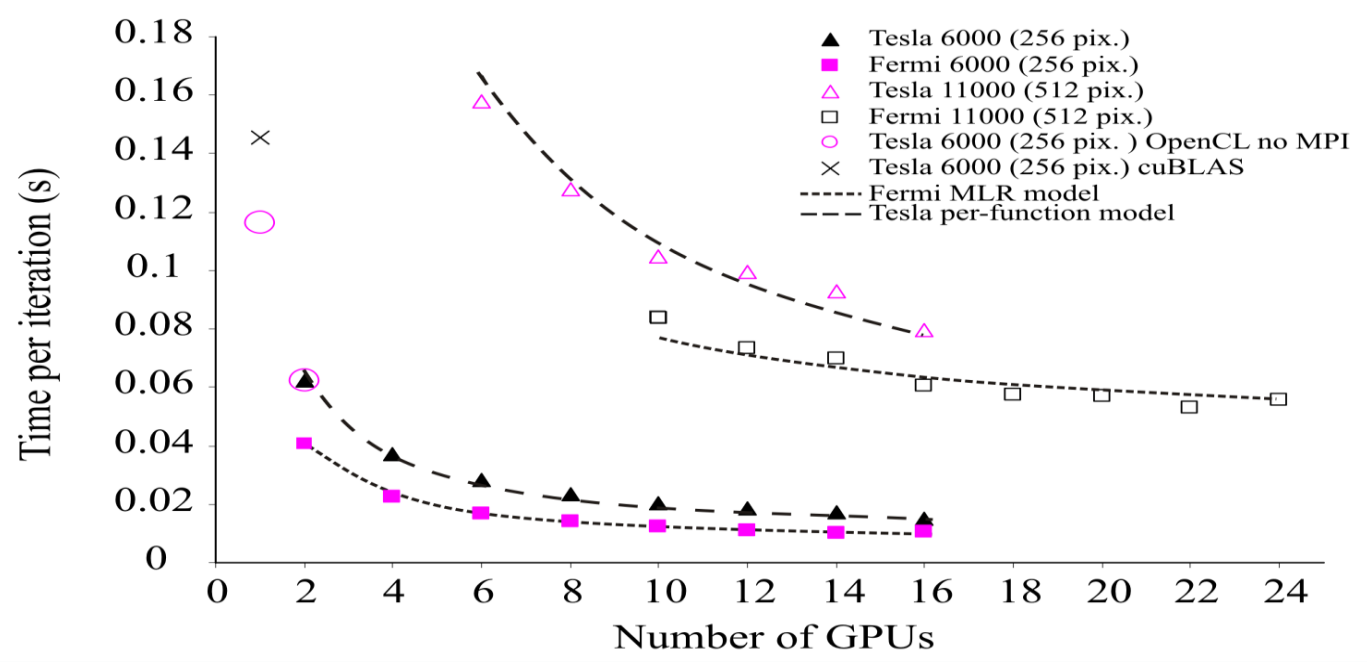

Instrument Design

"We've started to use the GPU cluster to speed up modelling of nuclear analysers such as CSIRO's air cargo scanner.
The speed is up to 5,000 to 10,000 times that of a normal desktop computer if we use most of the cluster. With this performance increase, simulations that normally take hours can be run interactively in real-time.
We expect this interactivity to significantly benefit the design and optimisation of new nuclear instruments."

Future of HPC

"The exciting part is that MIC and GPU foreshadow what will be on the CPU in the future. The work that scientists are putting in now to optimize codes for these processors will pay off. It's not whether you should adopt them; it's whether you want to get a jump on the future. "

Dan Stanzione, Texas Advanced Computing Center (TACC) funded by the National Science Foundation (NSF) to build Stampede.

Thank you

CSIRO Digital Productivity and Services Flagship

t +61 2 6216 7077

E John.A.Taylor@csiro.au

w www.csiro.au