Real-time visualisation and analysis of tera-scale datasets

Christopher Fluke

Amr Hassan (Swinburne; PhD student),
David Barnes (Monash University), Virginia Kilborn (Swinburne)

Thank you to the SPS15 organizers for the invitation to speak

Motivation

The Petascale Astronomy Data Era

MORE of the sky
MORE often
MORE pixels
MORE wavelengths
MORE data
MORE …

MORE computational work
MORE time passes before you can do…

MORE science

Desktop Astronomy

How long are YOU prepared to wait for an "interactive" response at your desktop?

Volume	Memory	Local disk
Gigascale	Yes	Yes
Terascale	No	Yes (slow)
Petascale	No	No

↓

Scalable Remote service

Australian SKA Pathfinder: Astronomy's Petascale Present

- 36 antennas
- Phased-array feeds
- Wide field of view
- 700 MHz – 1.8 GHz

2012-13: BETA ⟶ 2014: Full science

"*Hazards along the road include kangaroos, cattle, sheep, goats, goannas, eagles, emus, wild dogs….*"

http://www.atnf.csiro.au/observers/visit/guide_murchison.html#directions

Credit: Swinburne Astronomy Productions

WALLABY: The ASKAP HI All-Sky Survey

B.Koribalski (ATNF), L.Staveley-Smith (ICRAR) + 100 others…

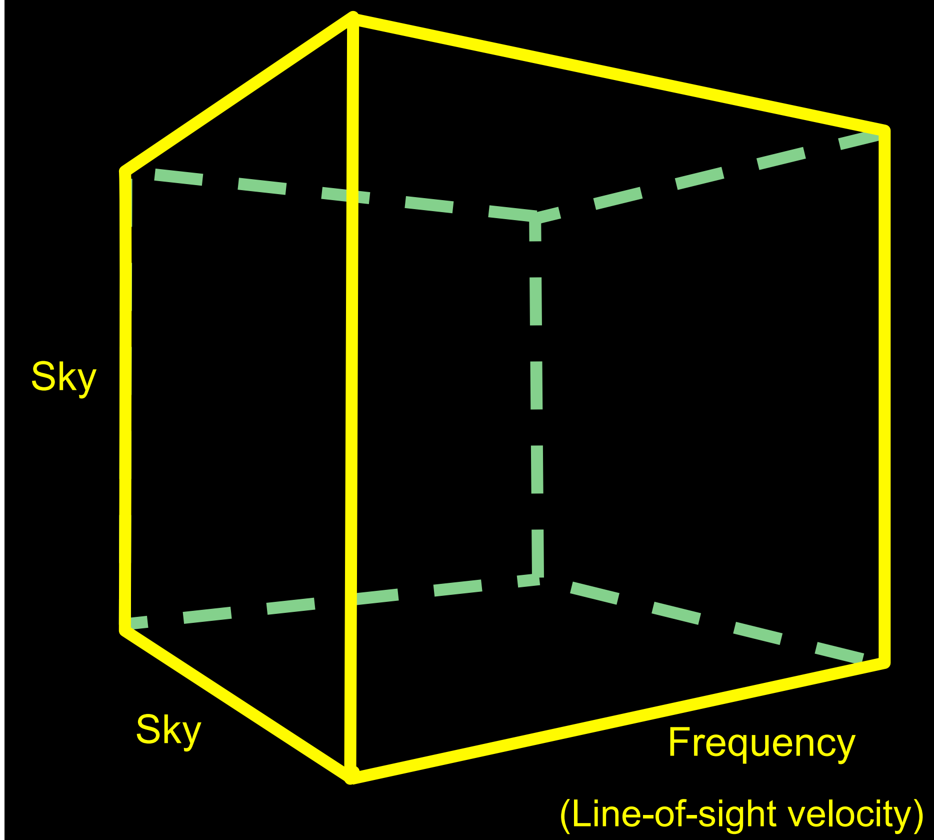

- Redshifted 21-cm HI
- ~0.5 million new galaxies
- 75% of sky covered
- z = 0.26 ~ 3 Gyr look-back

$$\nu = \sqrt{\frac{1 - v_{\rm los}/c}{1 + v_{\rm los}/c}}\,\nu_0$$

Observed

Emitted

Line-of-sight velocity

WALLABY: The ASKAP HI All-Sky Survey

B.Koribalski (ATNF), L.Staveley-Smith (ICRAR) + 100 others…

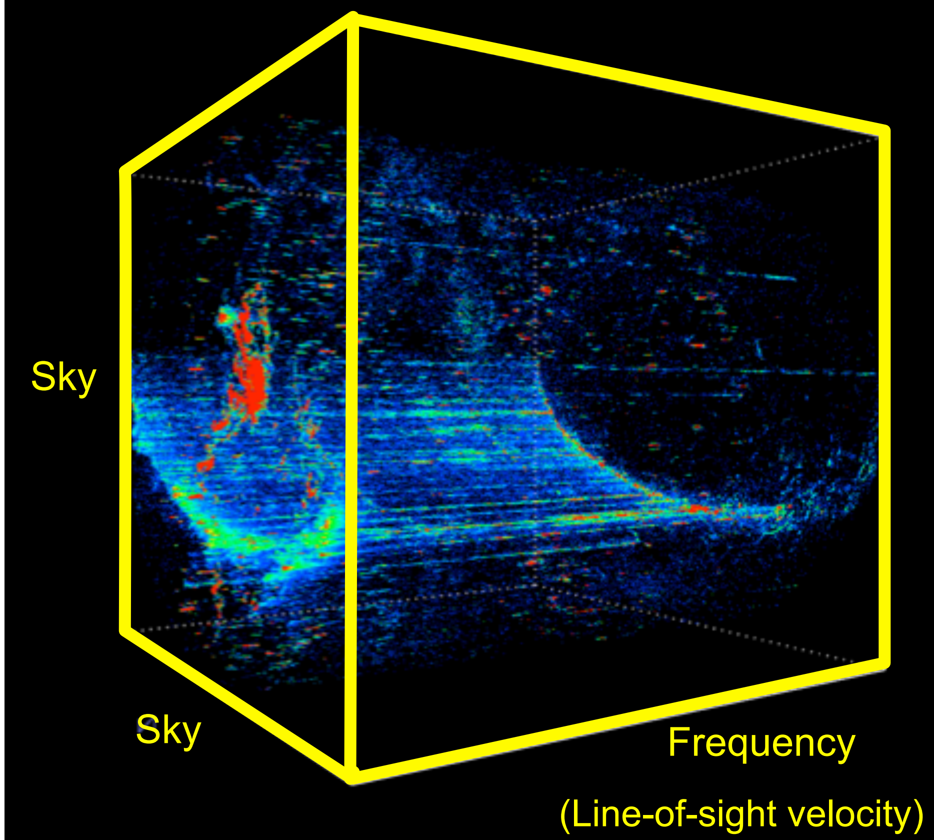

Likely data products:

4096 x 4096 x 16384 channels

~ 1TB per cube

[x1200 cubes]

Can we support real-time, interactive ***visualisation*** and ***data analysis?***

387 HIPASS cubes:
1721 x 1721 x 1024 = 12GB
Data: R. Jurek (HIPASS;ATNF)


gSTAR
GPU Supercomputer for Theoretical Astrophysics Research
sgi
Funding = AAL/Education Investment Fund + Swinburne
Peak: ~130 Tflop/s
100 x NVIDIA Tesla C2070 + 21 x NVIDIA Tesla M2090
Credit: Gin Tan

Graphics Processing Units (GPUs) are…

Massively parallel

Programmable*

Computational co-processors

Providing 10x-100x speed-ups

For many scientific problems

At low cost (TFLOP/$)

(But you can't use existing code)

[* CUDA, OpenCL, PyCUDA, Thrust, OpenACC, CUFFT, cuBLAS ….]

The future of computing is massively parallel

Run an individual problem faster

Save time

Lower price/ performance for Tflop/s HPC

Save money

Solve more complex problem in the same time

Increased accuracy

Solve bigger problem in the same time

Higher resolution

Run more problems in the same time

Parameter space

Is *my algorithm* suitable for a GPU?

See: Barsdell et al. MNRAS (2010), Fluke et al. PASA (2011)

Why types of problems are GPUs good for?

Inherent data parallelism

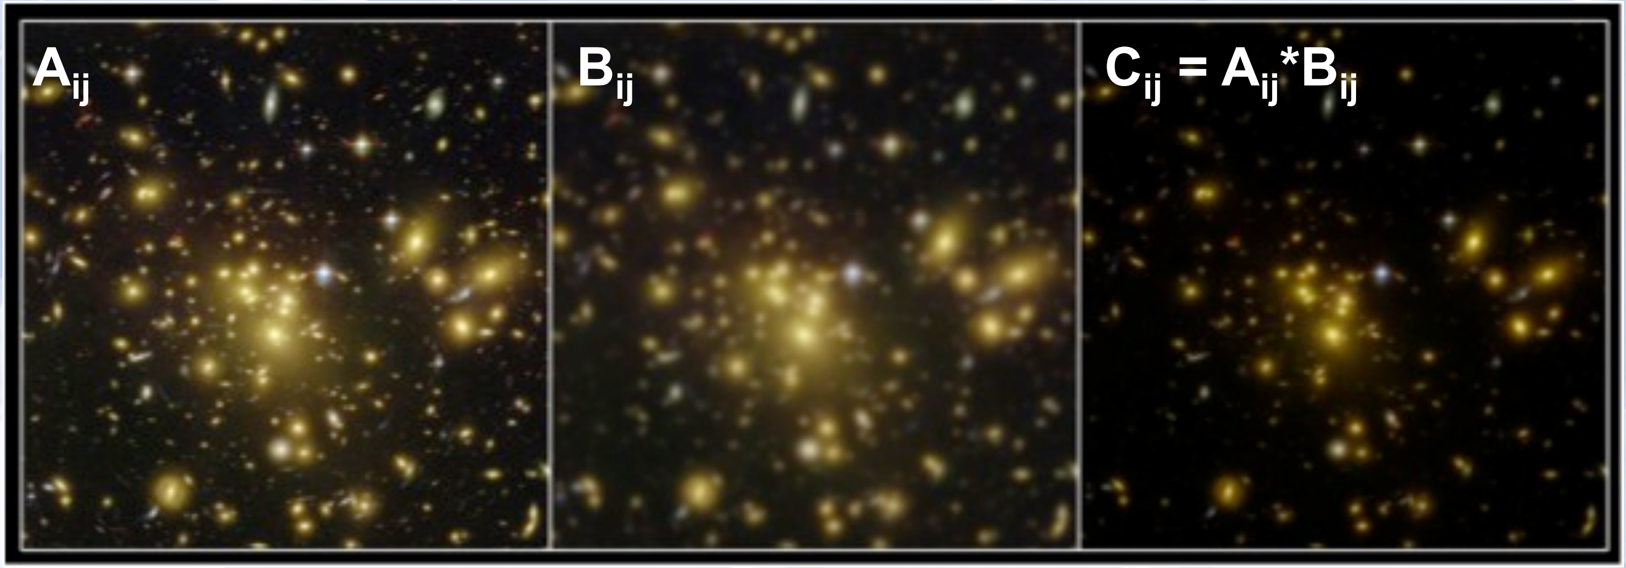

Abell 1689: NASA/Benitez et al.

E.g. pixel-by-pixel operations (SIMD)

High arithmetic intensity

$$\mathbf{y} = \begin{pmatrix} 1-\gamma & 0 \\ 0 & 1+\gamma \end{pmatrix} \mathbf{x} - \sigma_c \mathbf{x} - \sum_{i=1}^{N_*} m_i \frac{(\mathbf{x} - \mathbf{x}_i)}{|\mathbf{x} - \mathbf{x}_i|^2}$$

$N_* >> 1$

What are GPUs being used for in astronomy?

(ADS abstract search: 1 February 2012)

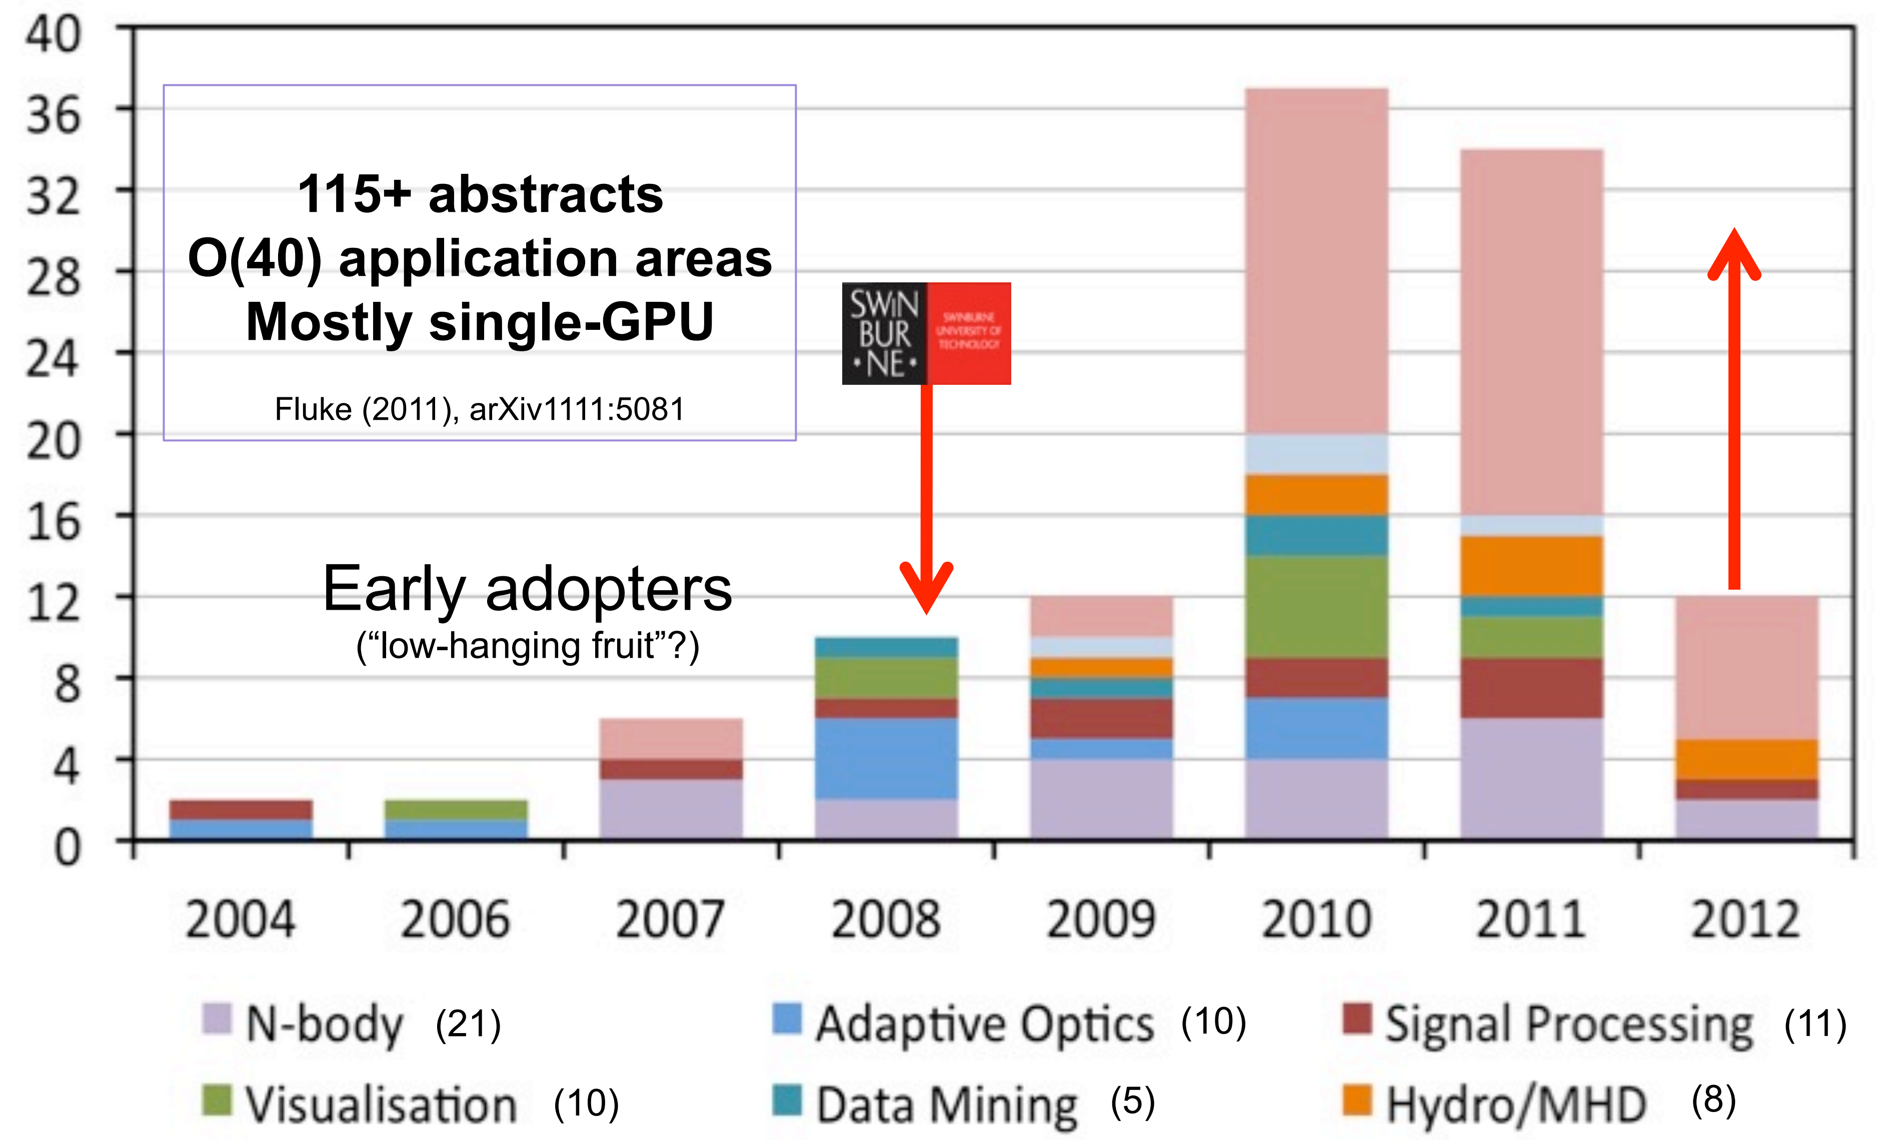

Volume Rendering via Ray Casting

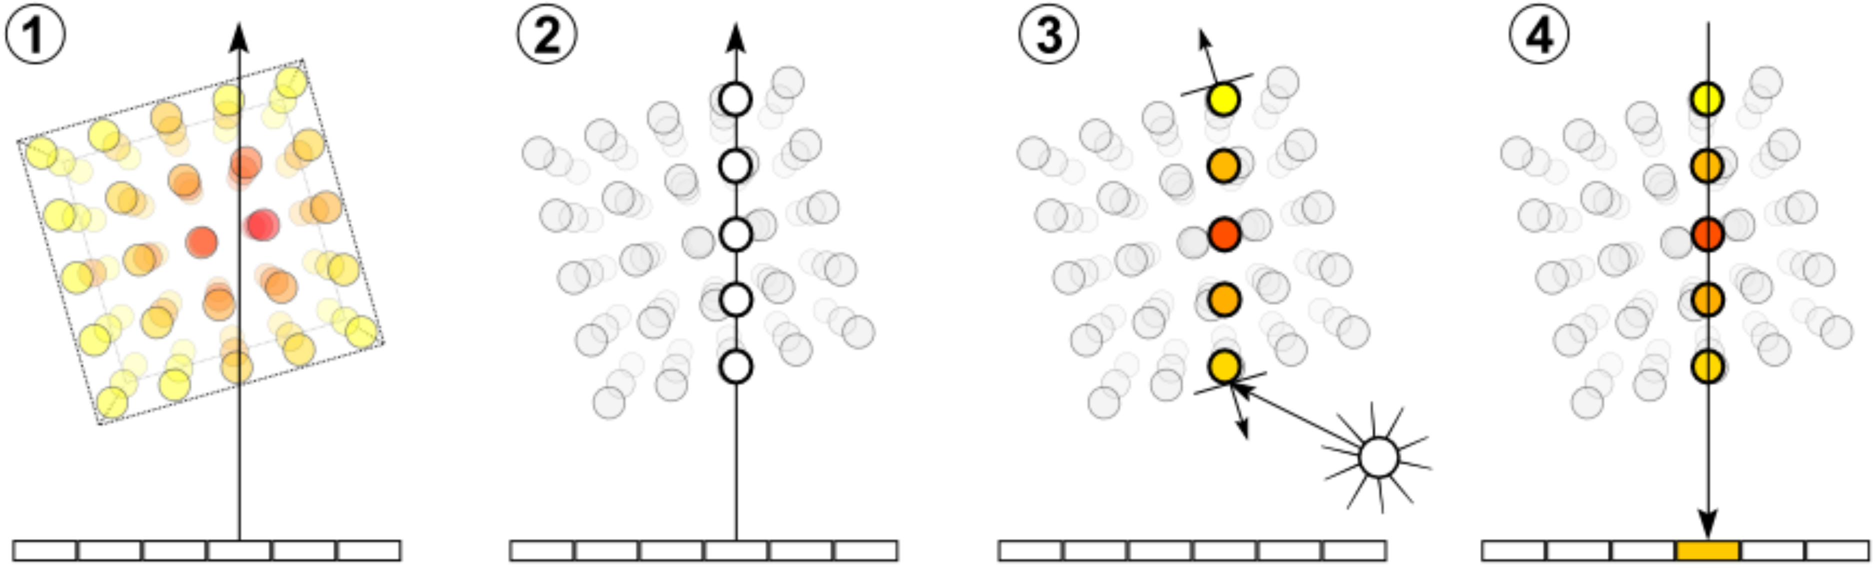

Ray casting

Sampling

Shading
Transfer function

Compositing

Data parallelism + high arithmetic intensity

Image: Wikimedia Commons

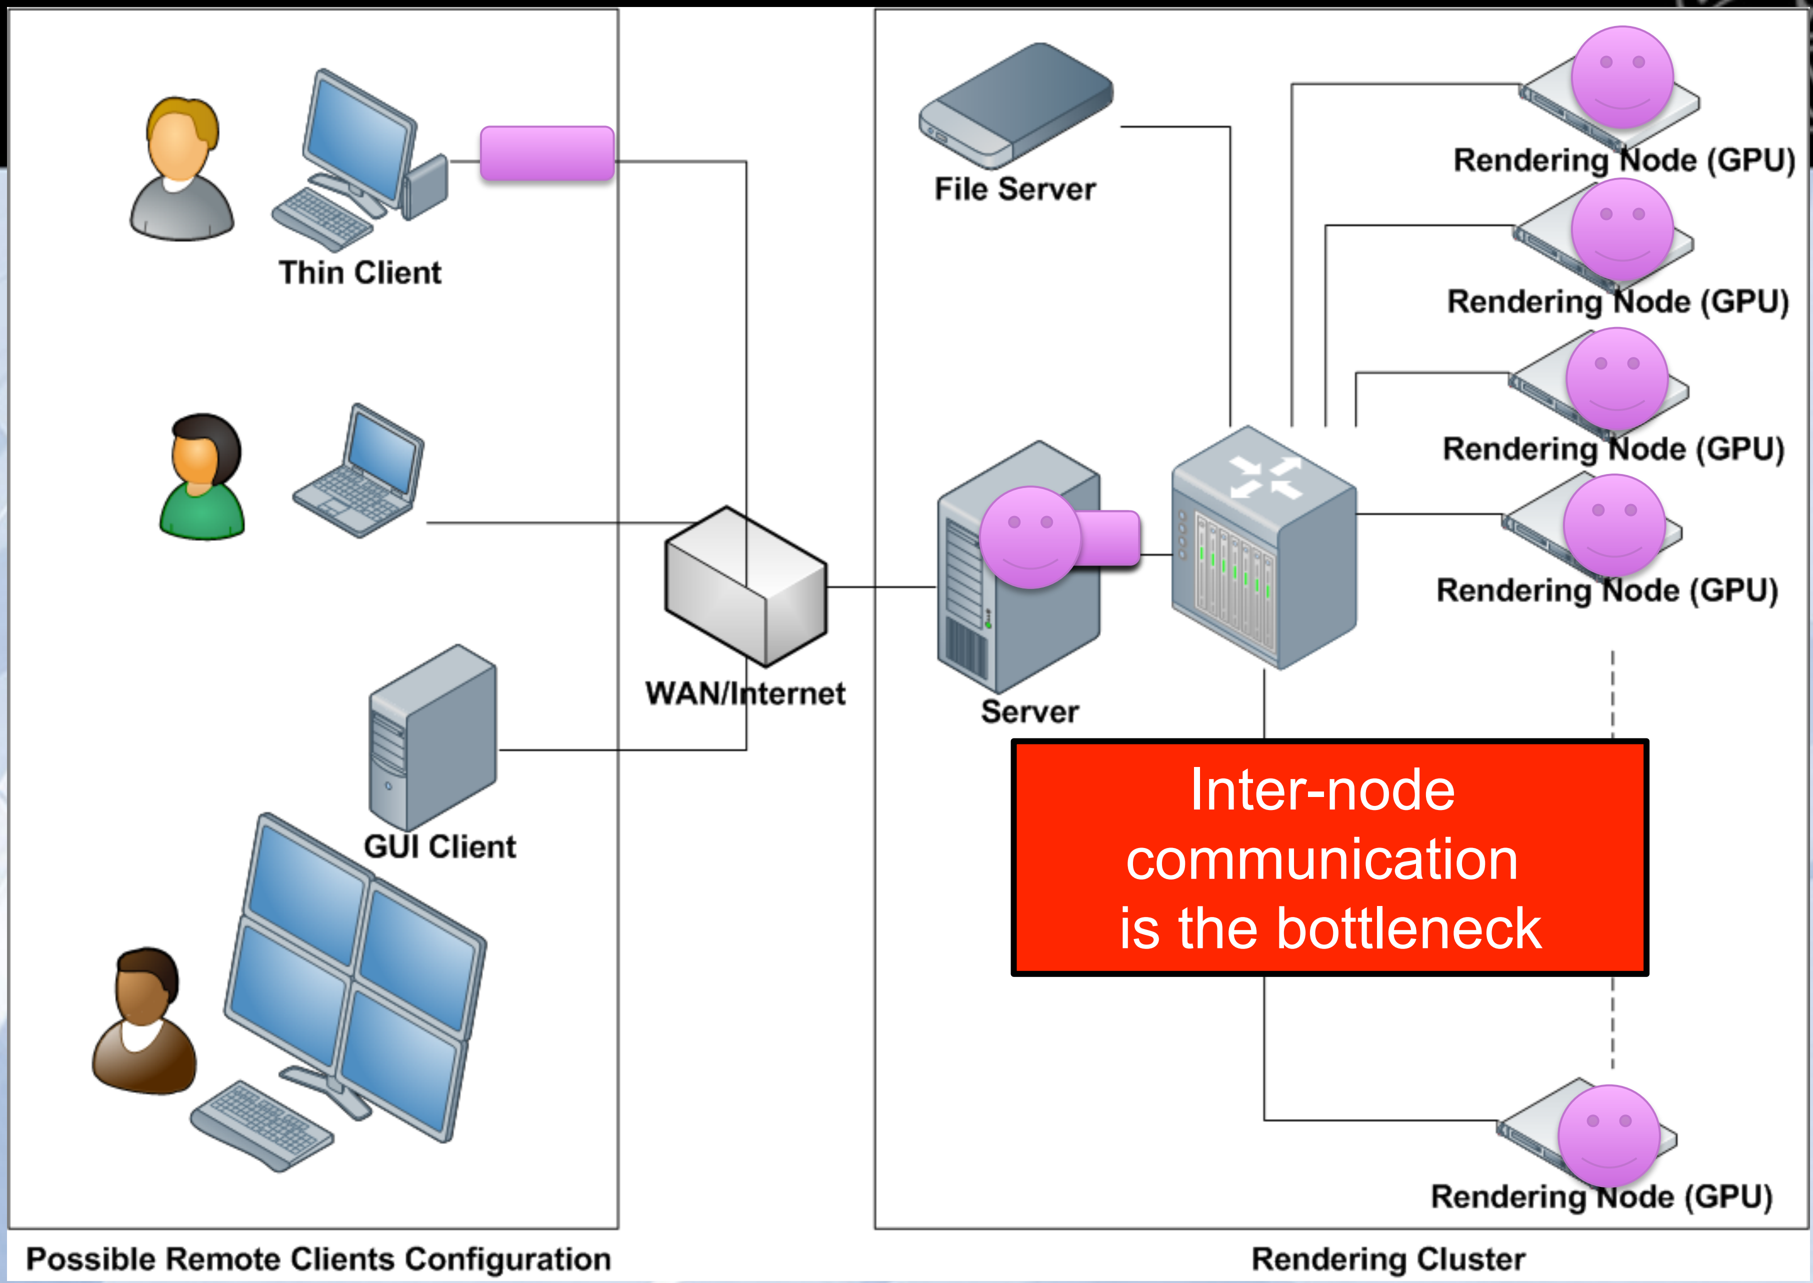

For details see: Hassan et al. (2010), NewA and Hassan et al. (2012), PASA

Early Benchmarking: Maximum Intensity Projection

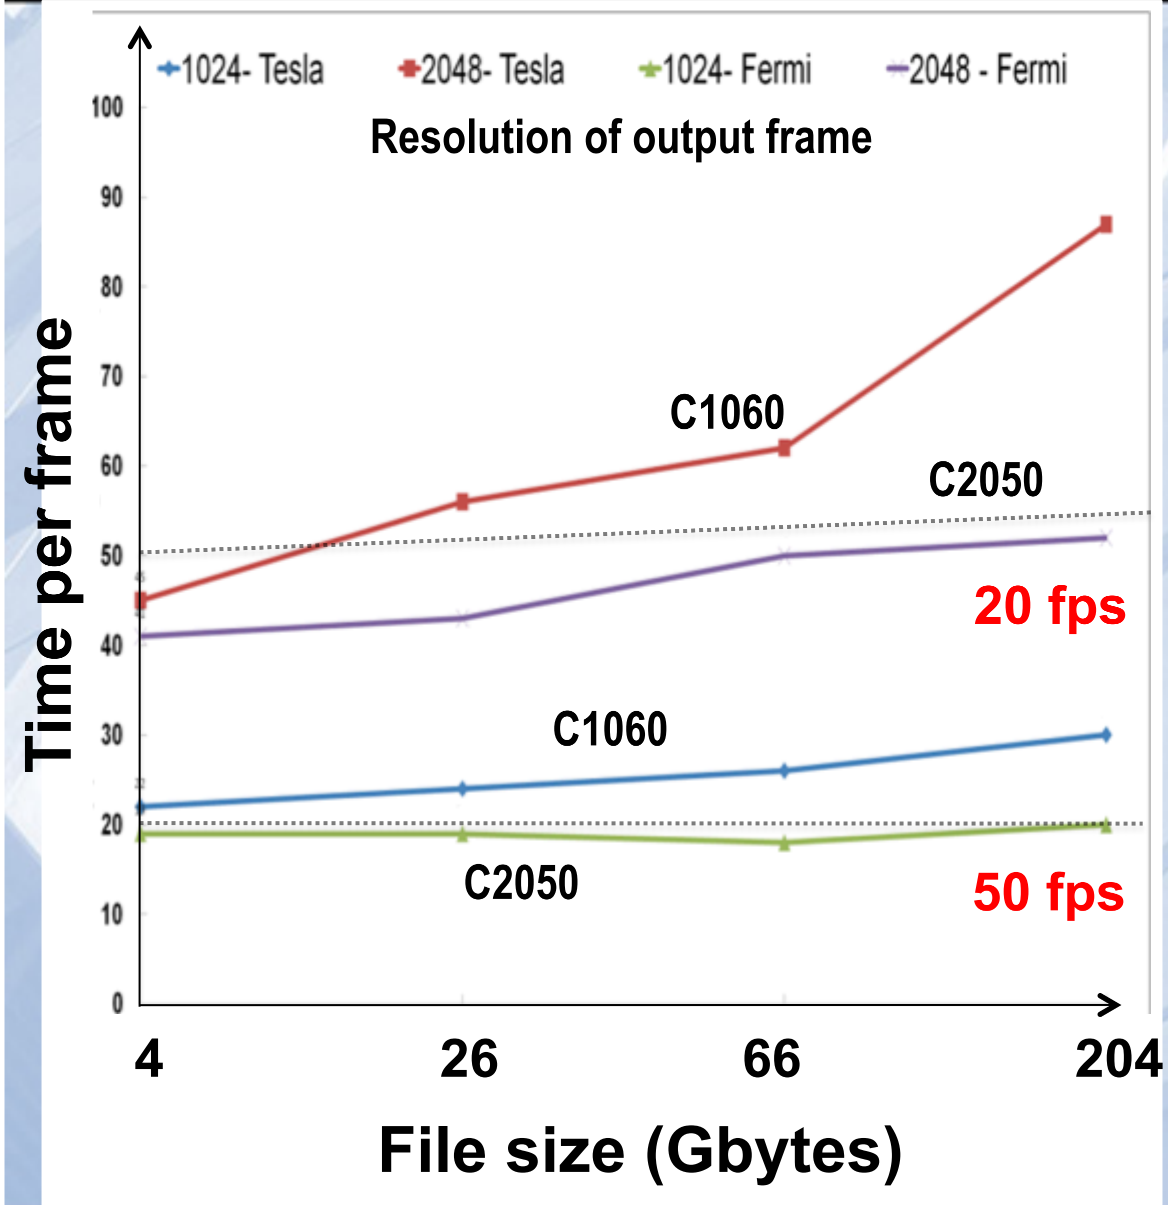

CSIRO GPU cluster

- 64 CPU nodes
- 128 GPUs
  - C1060 (older)
  - C2050 (newer)

Overhead = Inter-node communication

See Hassan et. al. (2012), PASA, online early

Framework enhancements (Hassan et al. 2012, submitted)

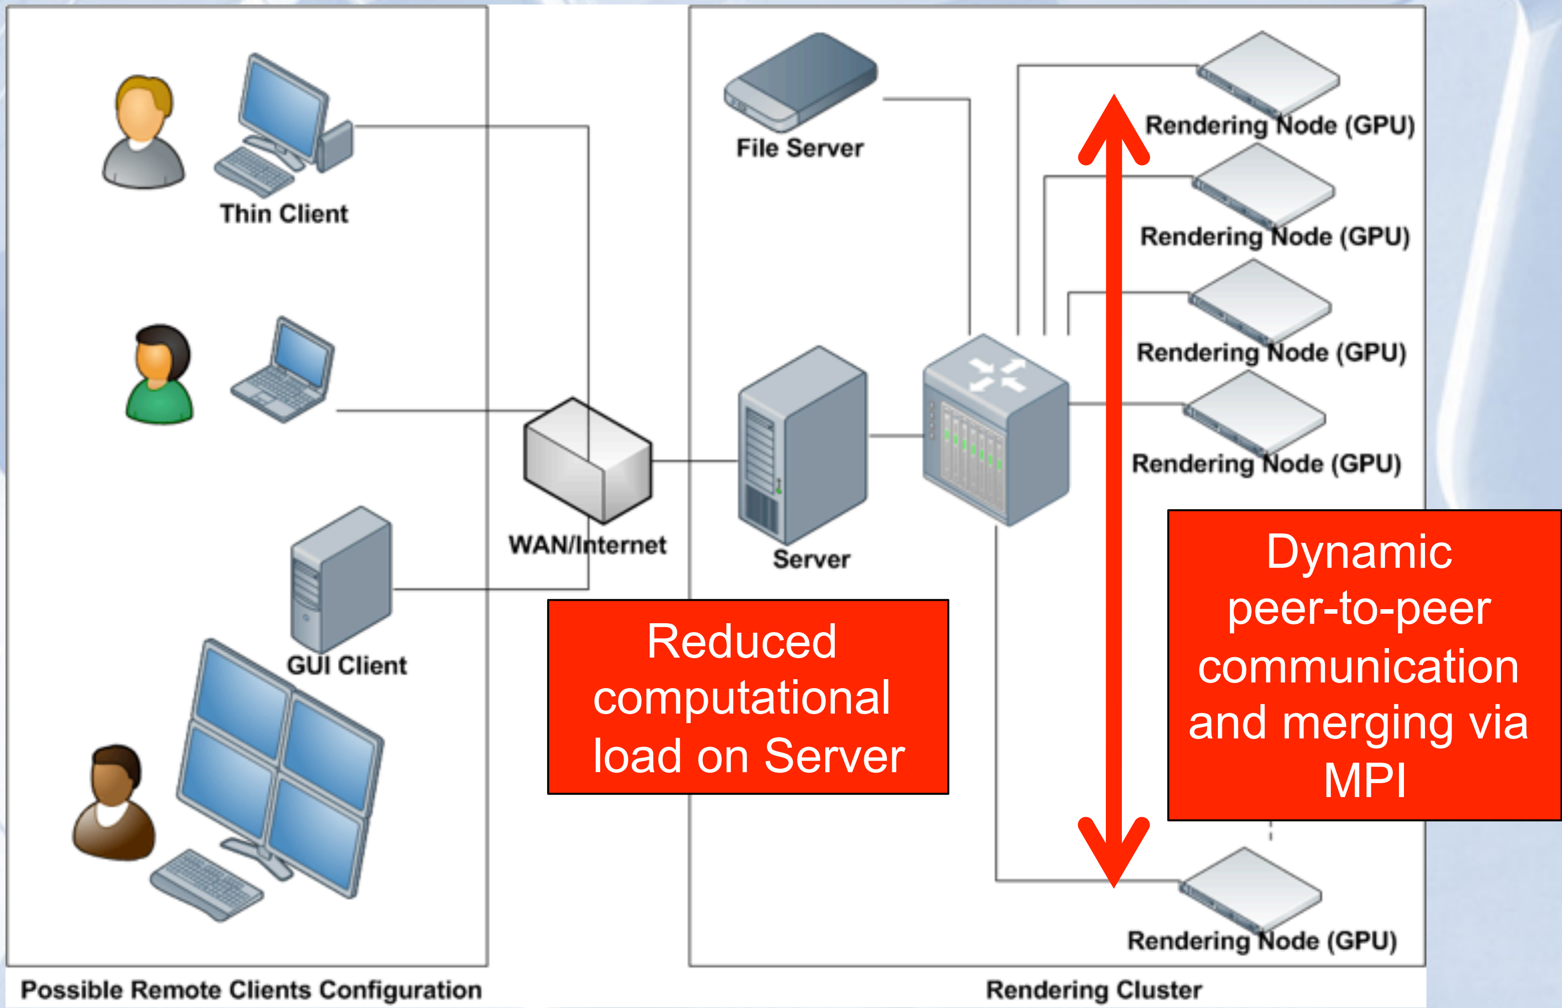

Supports arbitrary transfer function
= quantitative visualisation or data analysis

By the numbers: put the whole cube in memory

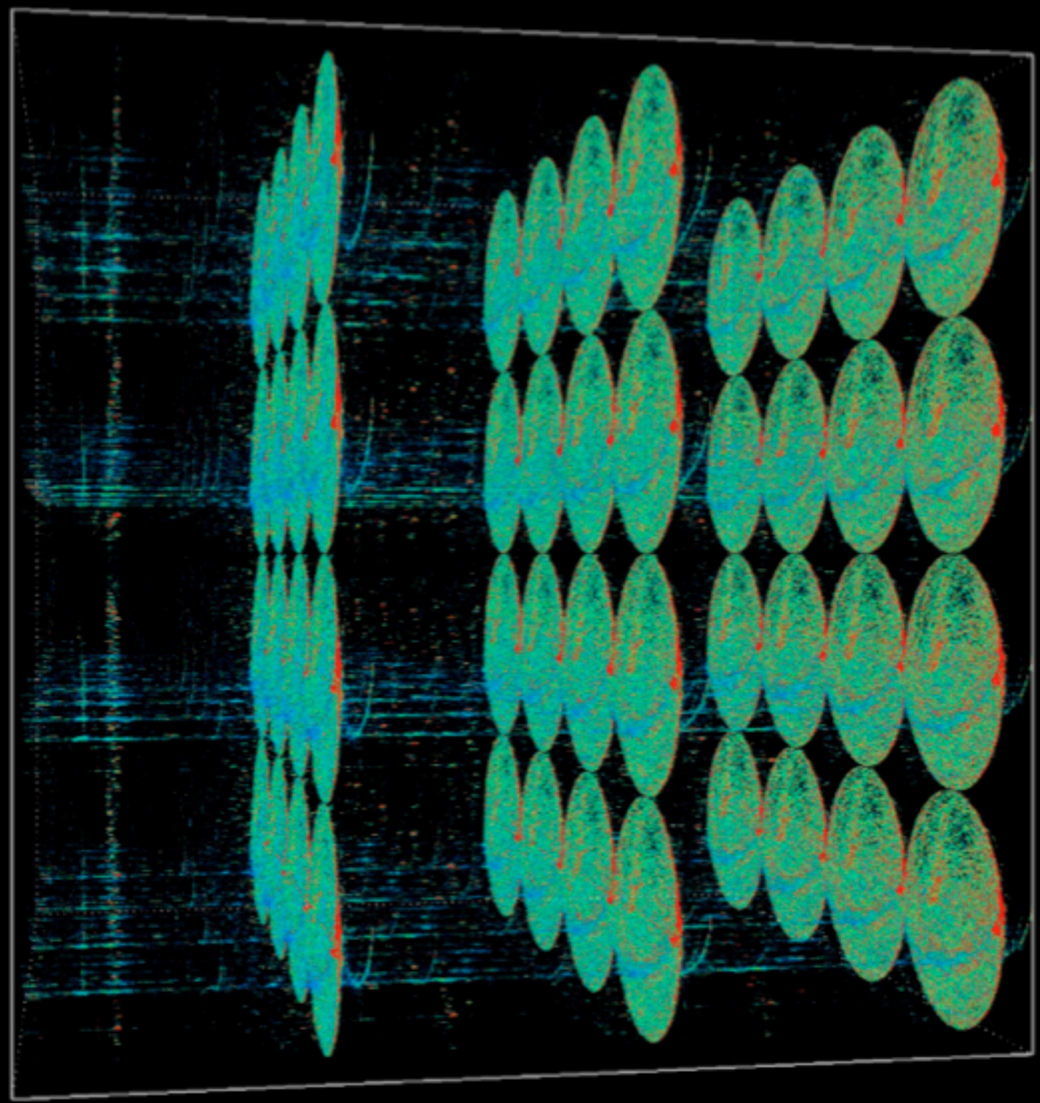

48 x HIPASS
- 4 x 4 x 3
- 6884 x 6884 x 3072
- 542.33 GB

96 GPUs
- 90 Tesla C2070
- 6 Tesla C2090
- 6 GB/GPU
- 43392 cores

Lustre file system
- 113 strips
- 546 sec = 9 min load

Visualisation: Scalability Testing

Configuration	Facility	Maximum size	Tested
32 node – 64 GPU (3GB/GPU) Minimum 128 CPU cores	CSIRO GPU Cluster	140 GB	Yes
64 node – 128 GPU (3GB/GPU) Minimum 256 CPU cores	CSIRO GPU Cluster	281 GB	Yes
32 node – 64 GPU (6GB/GPU) Minimum 128 CPU cores	**gSTAR**	**300 GB**	**Yes**
48 nodes – 96 GPU (6GB/GPU) Minimum 192 CPU cores	**gSTAR**	**540 GB**	**Yes**
64 nodes – 128 GPU (6GB/GPU) Minimum 256 CPU cores	Upgrade (2012?)	650 GB	Planned
128 nodes – 256 GPU (6GB/GPU) Minimum 512 CPU cores	Upgrade (2013?)	1.3 TB	No

> 10 fps (first three rows)

~ 7fps (48 nodes row)

WALLABY: 2014!

Analysing **0.5 Tbyte** (on 96 GPUs)

Task	Description	Time
Histogram	Visit each data point once	~4 sec
Global mean and standard deviation	Summarizing whole dataset into single value(s)	~2 sec
Global median	Multiple iterations to convergence (Torben's method)	~45 sec
3D spectrum tool	Quantitative data interaction: click for spectrum	20 msec

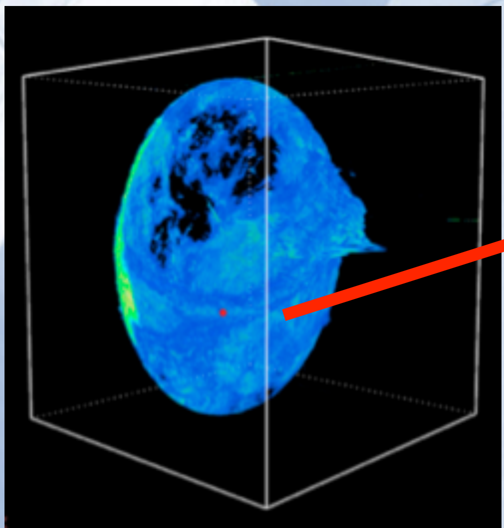

Interactive 3D quantitative visualisation

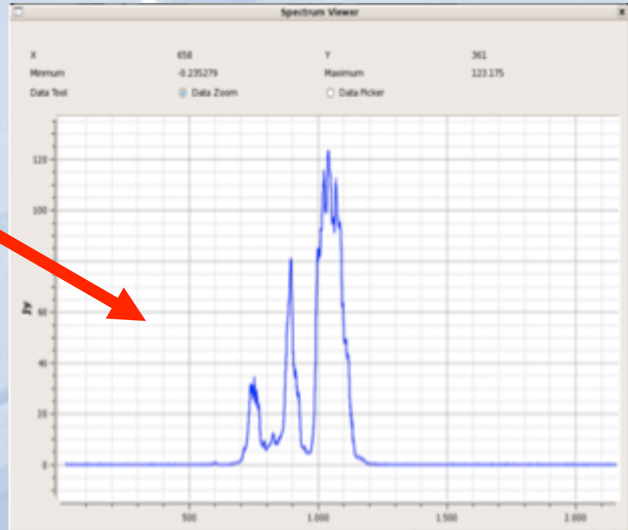

Data: GASS (N.McClure-Griffiths; ATNF)

Interactive data thresholding

2σ

3σ

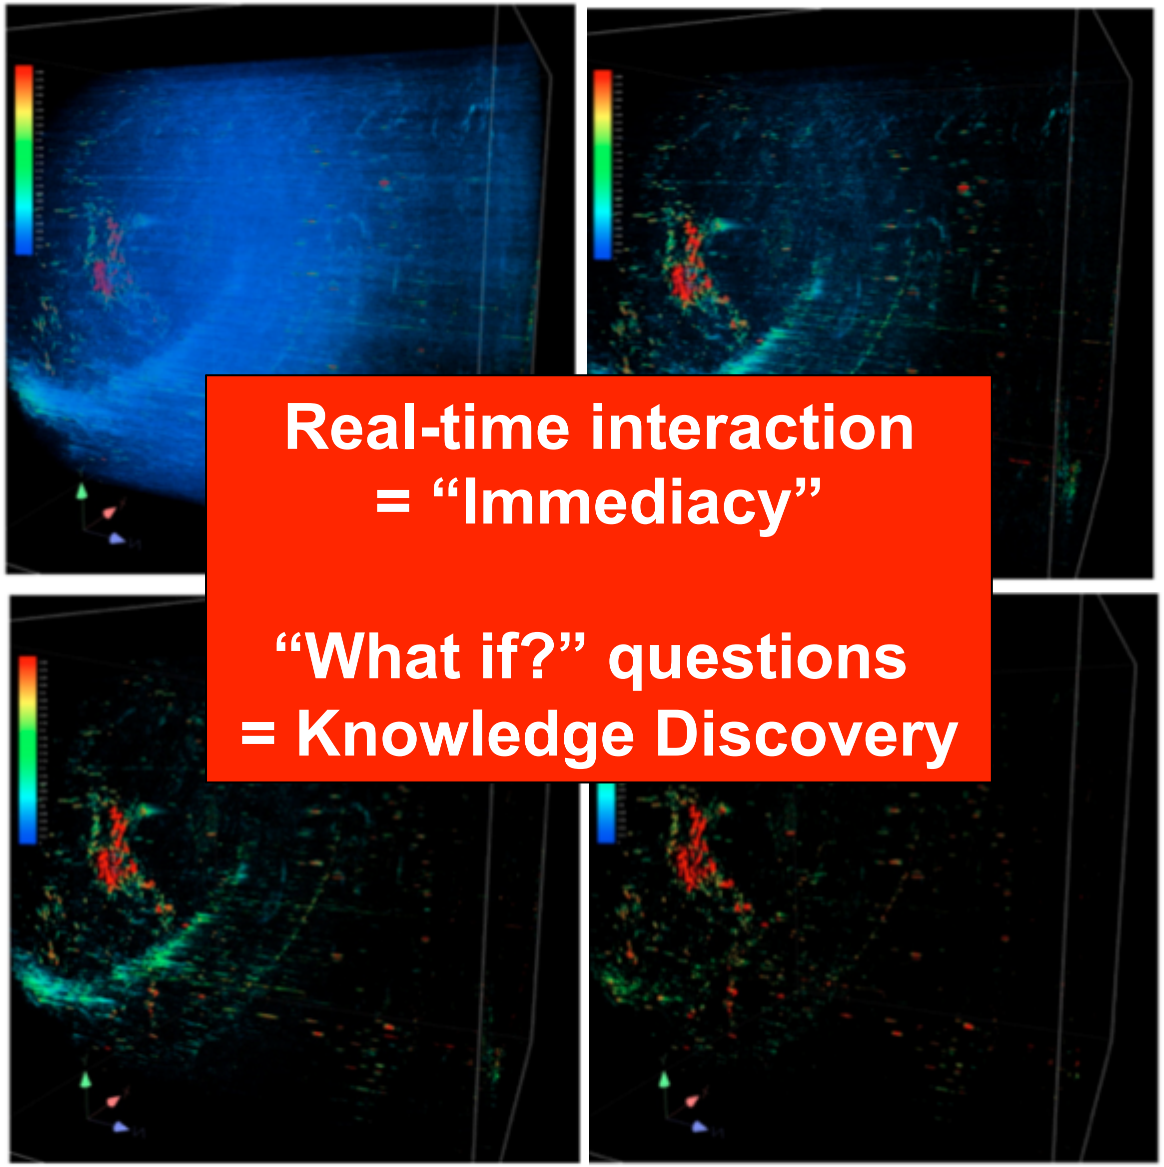

Real-time interaction
= “Immediacy”

“What if?” questions
= Knowledge Discovery

4σ

7σ

Hassan et al.
2012, submitted

Future directions?

- Large-format displays
- Temporal data
- Polarisation (Stokes)
- New transfer functions
  - E.g. medical imaging

8000×8000 pixel volume rendering of the HIPASS dataset on the CSIRO Optiportal at Marsfield, NSW. Data: R. Jurek (ATNF) from 387 HIPASS cubes.
Image: C.Fluke

Conclusions

- Terascale real-time, interactive visualisation and data analysis?
  - **Achievable with GPU clusters**
  - Communication bound
- Wish list
  - More memory/GPU
  - More GPU/node (PCIe limit)
  - Faster inter-node communication
- Exciting parallel future!