Accelerating the rate of astronomical discovery with GPU-enabled clusters

Dr Christopher Fluke
Scientific Computing & Visualisation Group

ADASS 2011

Thanks to B.Barsdell (Swin), A.Hassan (Swin), D.Barnes (Monash) and ADASS POC

Computation in Astronomy

Wikimedia Commons

CJF

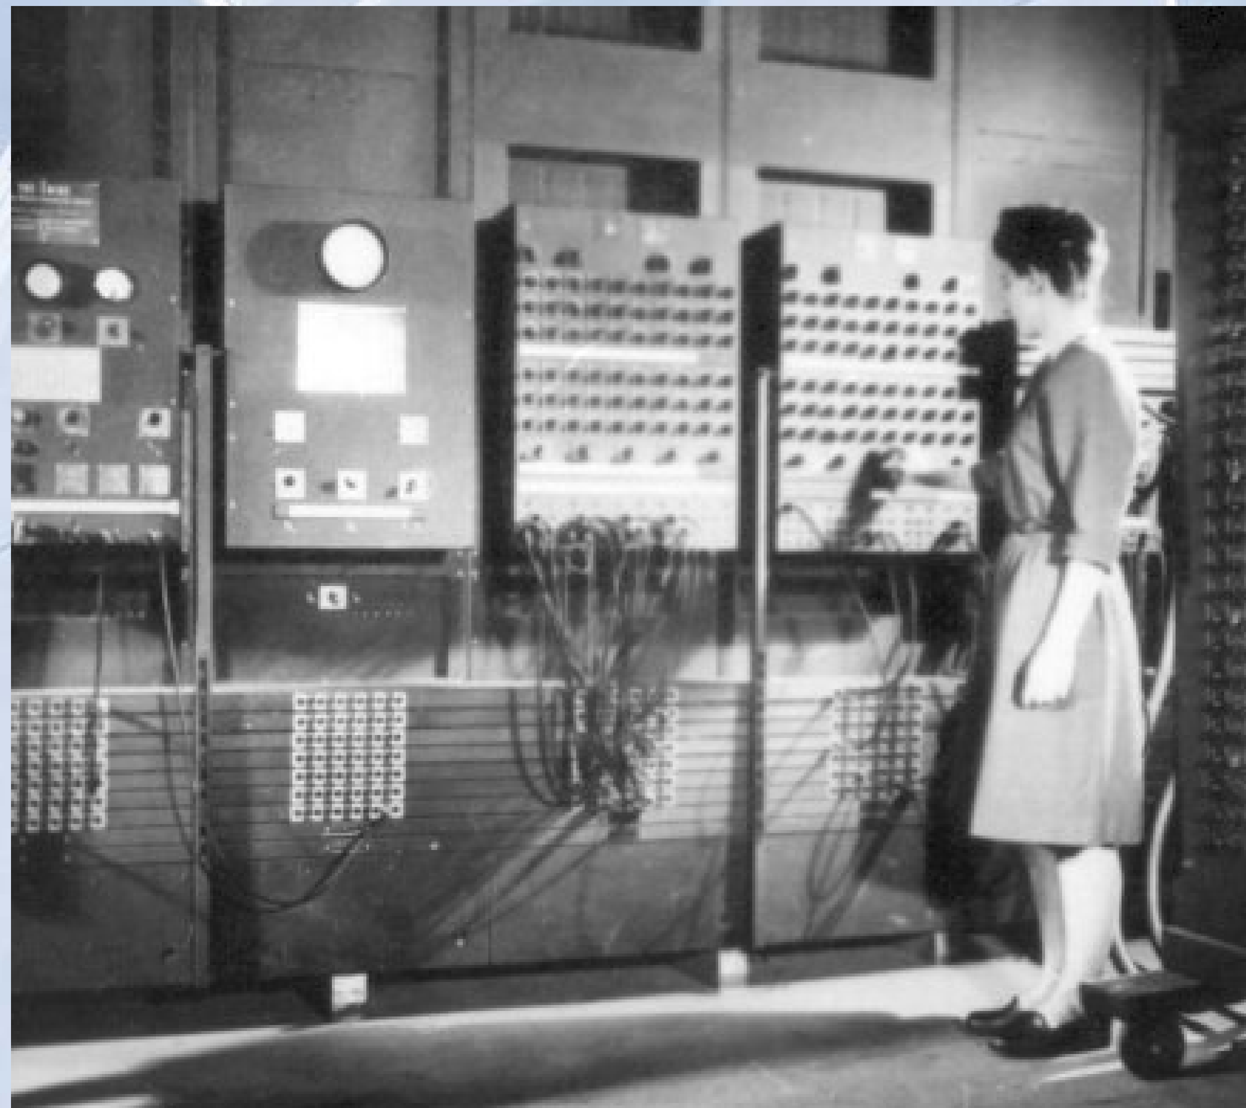

U.S. Army Photo, Wikimedia

This…

http://archive.gamespy.com/legacy/halloffame/hof-spaceinvaders/spaceinvaders3.gif

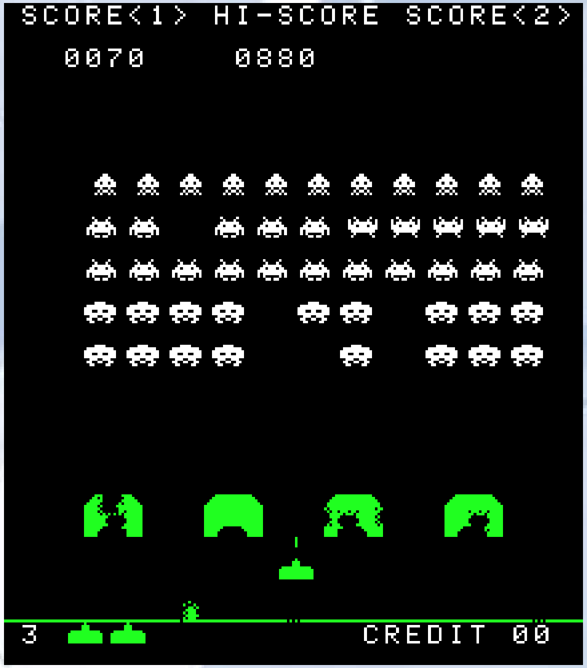

Now looks like this…

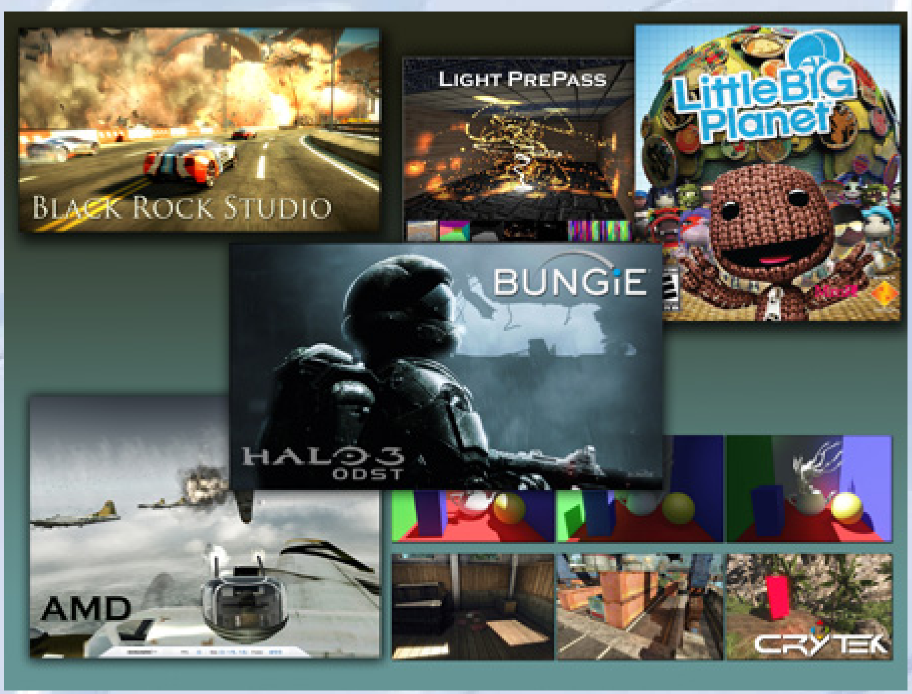

http://www.bungie.net/News/content.aspx?link=Siggraph_09

Thanks to devices like these…

Images: Wikimedia commons

Graphics Processing Units (GPUs)

- Programmable computational co-processor
- Low-cost “desktop supercomputer”
  - Offers better FLOP/$
  - Offers better FLOP/W
- Offer 10x-100x speed-ups for many science problems

NVIDIA
Tesla
C2075

1.03 TFLOP/s
(sp)515 GFLOP/s
(dp)

Image: http://www.nvidia.com

AMD Firestream
9350

2.64 TFLOP/s (sp)
528 GFLOP/s (dp)
2.4 GFLOPS/W

Image: http://www.amd.com

Motivation: Moore's Law

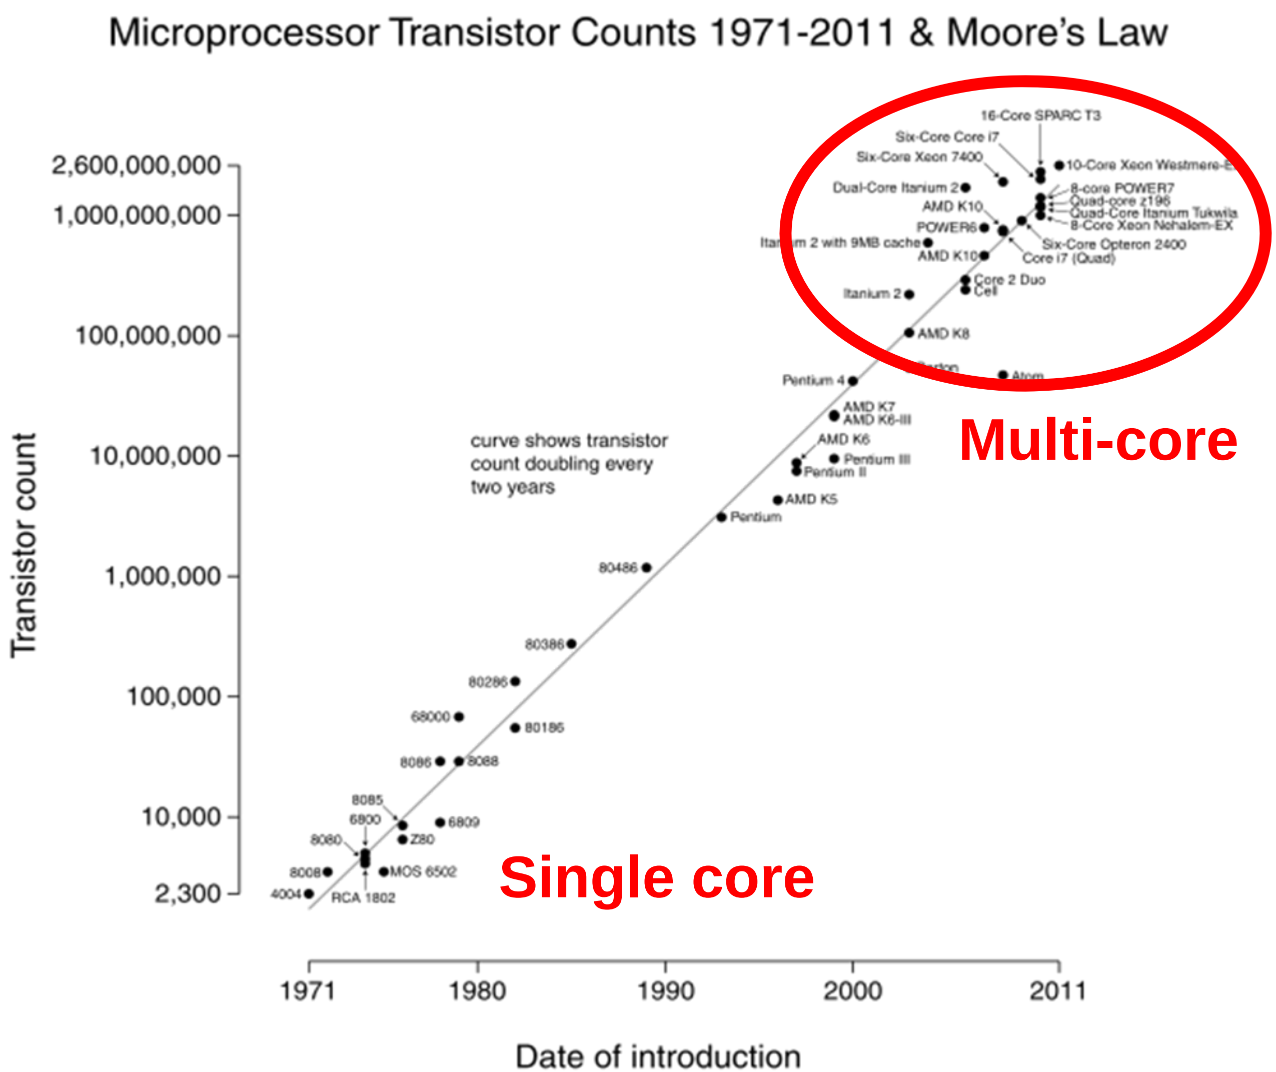

Image: Wikimedia commons

Motivation: The Multi-Core Corner

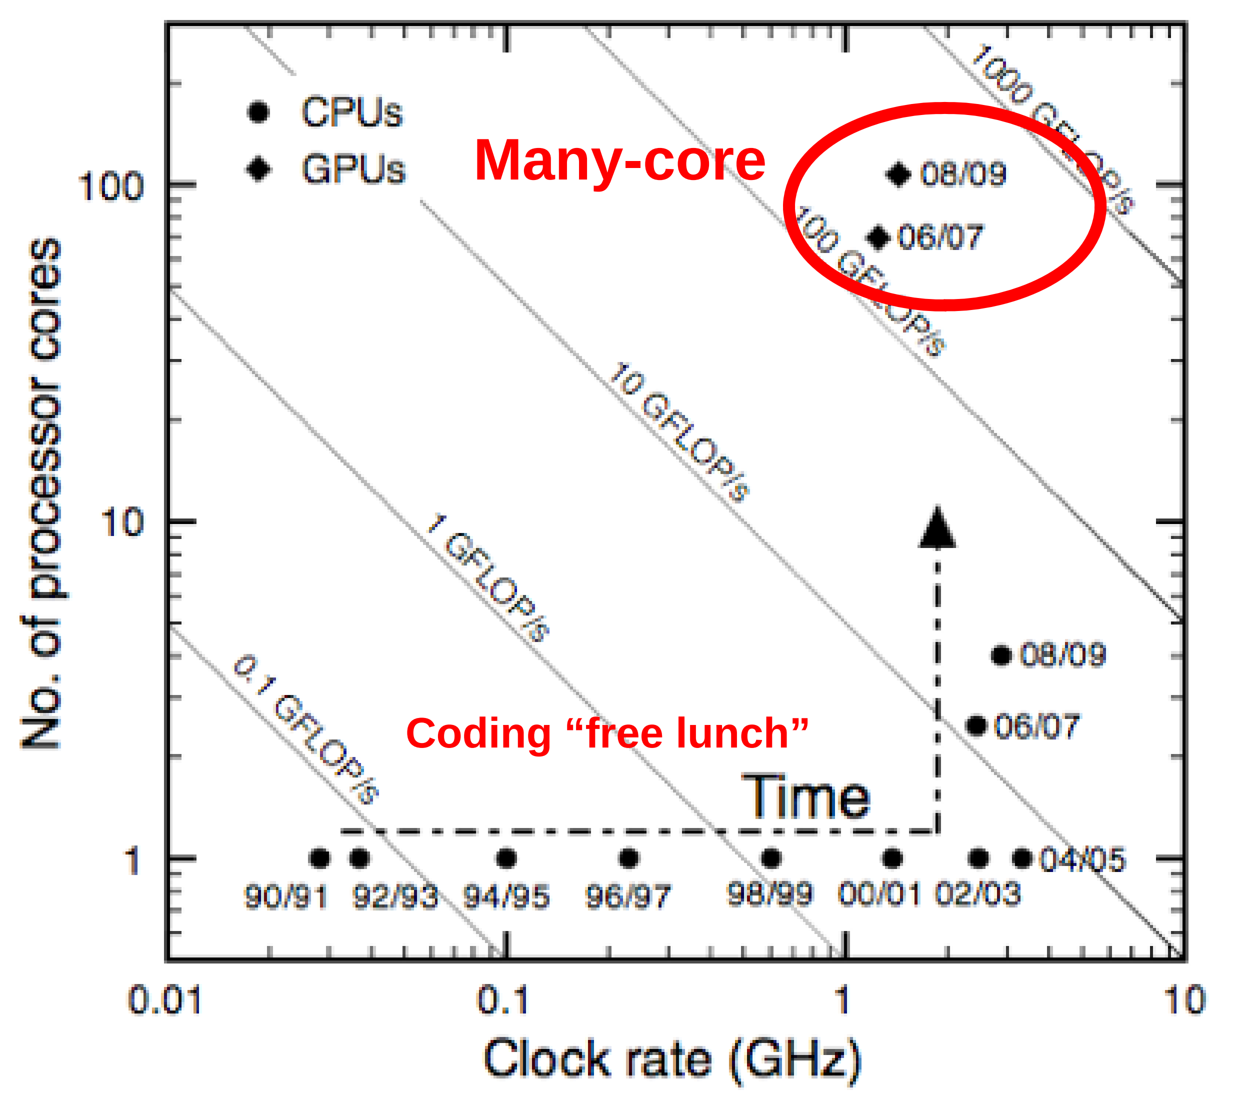

Image: B.Barsdell

CPUs vs. GPUs

CPUs:

- Have large-memory caches, sophisticated control logic
- Because they have to do everything
- They are relatively easy to program for any task

GPUs:

- Have circuit area devoted to floating point computations
- They are somewhat harder to program
- Because they were designed to do graphics
- “Single instruction multiple data” (SIMD)

GPUs for Scientific Computation

- General Purpose computing on GPU (GPGPU)

- Programmable pipeline
  - Shader languages: Cg; OpenGL; …
  - Application Programming Interfaces (APIs):
    - CUDA (NVIDIA – http://www.nvidia.com/cuda)
    - OpenCL (Khronos – http://www.khronos.org/opencl)
  - Growing number of other options
    - Thrust, PyCuda, ...

Early Adoption in Astronomy

N-body forces:

- $O(N^2)$ = High arithmetic intensity!
- Nyland, Harris, Prins (2004); NVIDIA GPU using Cg/OpenGL
- Elsen et al. (2006; 2007); ATI GPU using BrookGPU
- 20x speed-up compared to CPU
- Performance comparable to custom GRAPE-6A

Adaptive optics wave-front reconstruction

- Rosa et al. (2004)
- Recovery of wave-front phase from Shack-Hartmann sensor
- 10x speed-up for centroid calculation
- 2x speed-up overall

Early Adoption in Astronomy

Common-Off-the-Shelf (COTS) Correlator

- Schaaf & Overeem (2004)
- NVIDIA GeForce 6800 Ultra GPU vs. 2.8 GHz CPU
- ~5x better performance for 16x bigger problem
- Price/Gflop and Power/Gflop were 3x better for GPU

Emerging Trends (Amateur-ish Bibliometrics)

- ADS Abstract search
  - *GPU(s), graphics processing unit(s), CUDA, OpenCL*

- 94 abstracts…however…
  - Fails to find papers that use GPUs but don't have in abstract
  - Fails to find papers that use GPUs for astro but not in ADS

- Summary:
  - 3 classes (methods, science result, philosophy)
  - 30 broad application areas
  - ~50 unique computational problems

Classification

Application Area	2004	2006	2007	2008	2009	2010	2011	Totals
N-body (direct/tree/symplectic integration)			3	2	4	4	5	**18**
Adaptive Optics	1	1		4	1	3		**10**
Visualisation		1		2		5	2	**10**
Signal processing	1		1	1	2	2	2	**9**
Hydro/MHD					1	2	4	**7**
Adaptive Mesh Refinement					1	2	1	**4**
Data mining				1	1	2		**4**
Philosophy						3		**3**
3D Radiative Transfer						1	2	**3**
Microlensing						2		**2**
Spectral energy distribution						1	1	**2**
Spherical Harmonic transforms						1	1	**2**
Binary black hole inspirals						2		**2**
Cherenkov Radiation						1		**1**
Cosmological lattice						1		**1**
Galaxy fitting							1	**1**
Lomb-Scargle periodogram						1		**1**
LSST Source catalog simulations						1		**1**
Ly-Alpha forest simulations							1	**1**
Protoplanetary disks						1		**1**
Radial convolution kernels; HEALPix							1	**1**
Radiation belt simultions			1					**1**
Real-time, fast transients							1	**1**
Solving Kepler's equations					1			**1**
Two-point correlation function (FPGA)			1					**1**
Celestial object classification							1	**1**
Thermal planet models							1	**1**
Numerical Relativity						1		**1**
Gravitational waves						2		**2**
Software framework					1			**1**
Totals	**2**	**2**	**6**	**10**	**12**	**38**	**24**	**94**

Methods (82)
Science results (9)
Philosophy (3)

What are GPUs being used for? (1 October 2011)

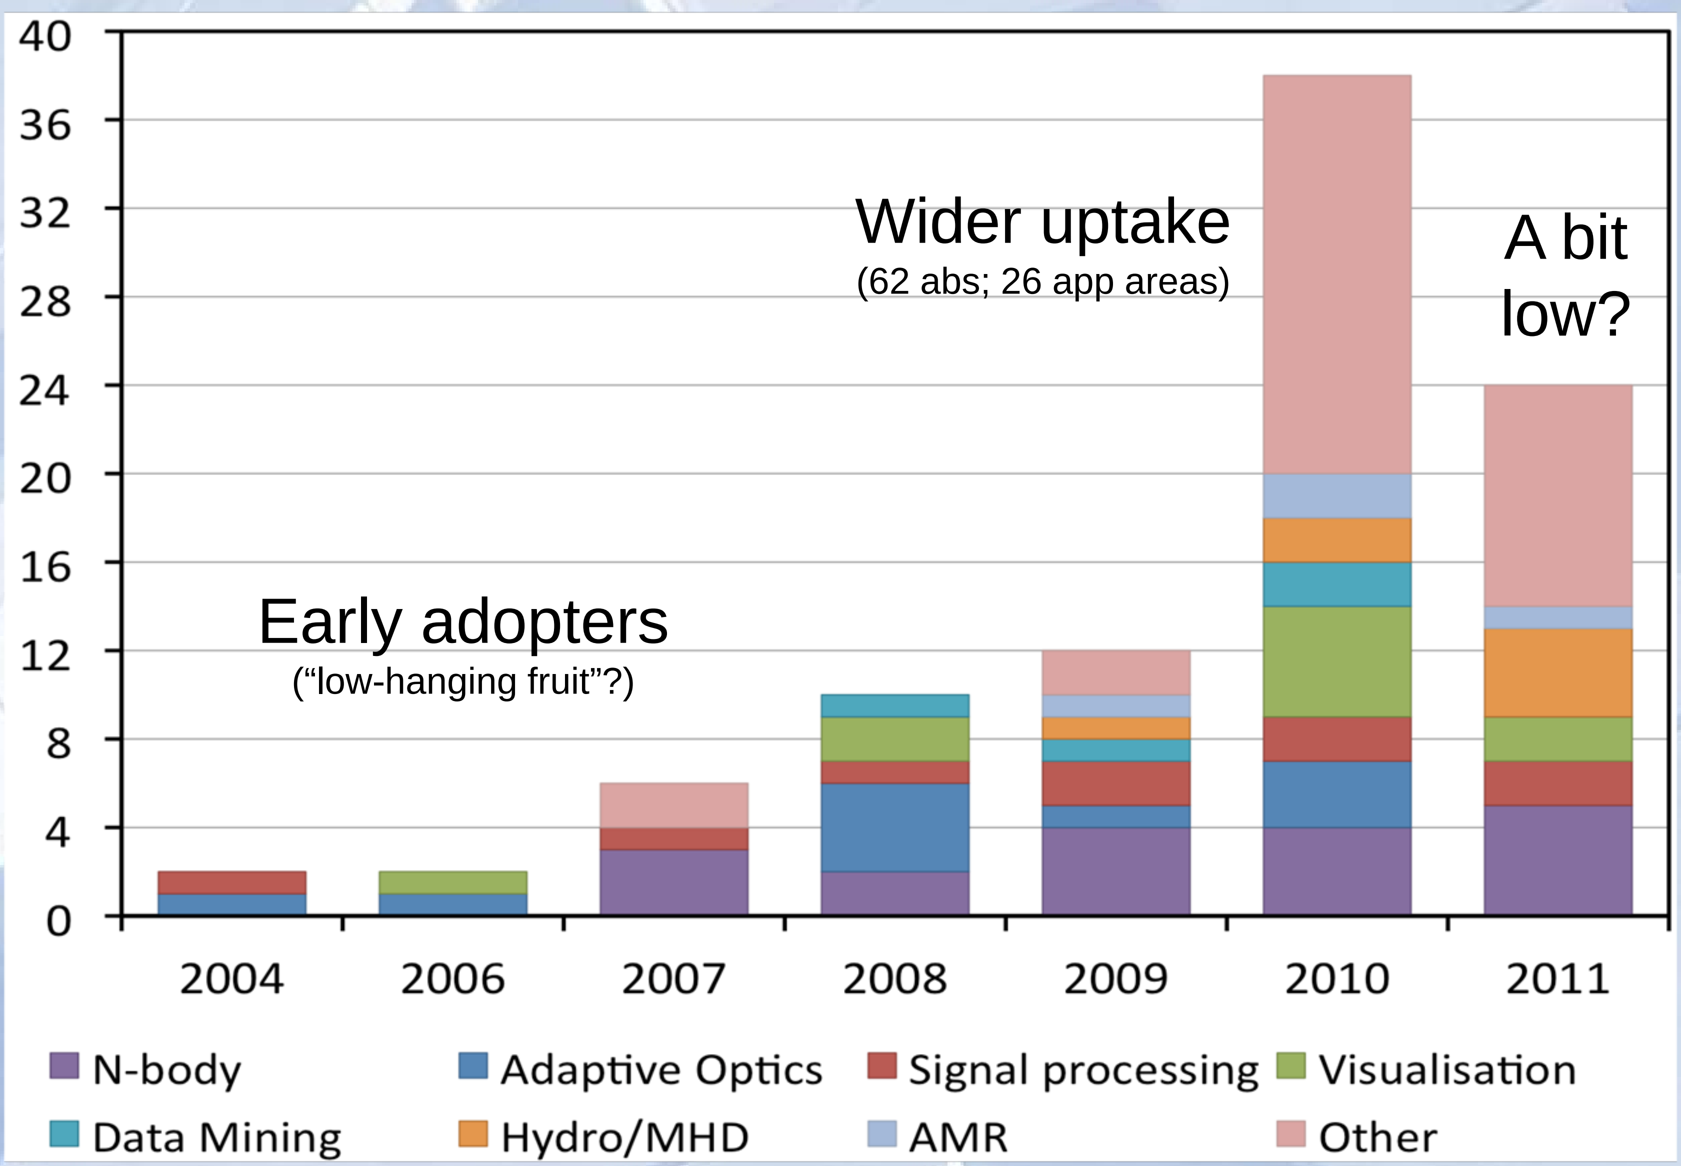

Where is it being published? (1 October 2011)

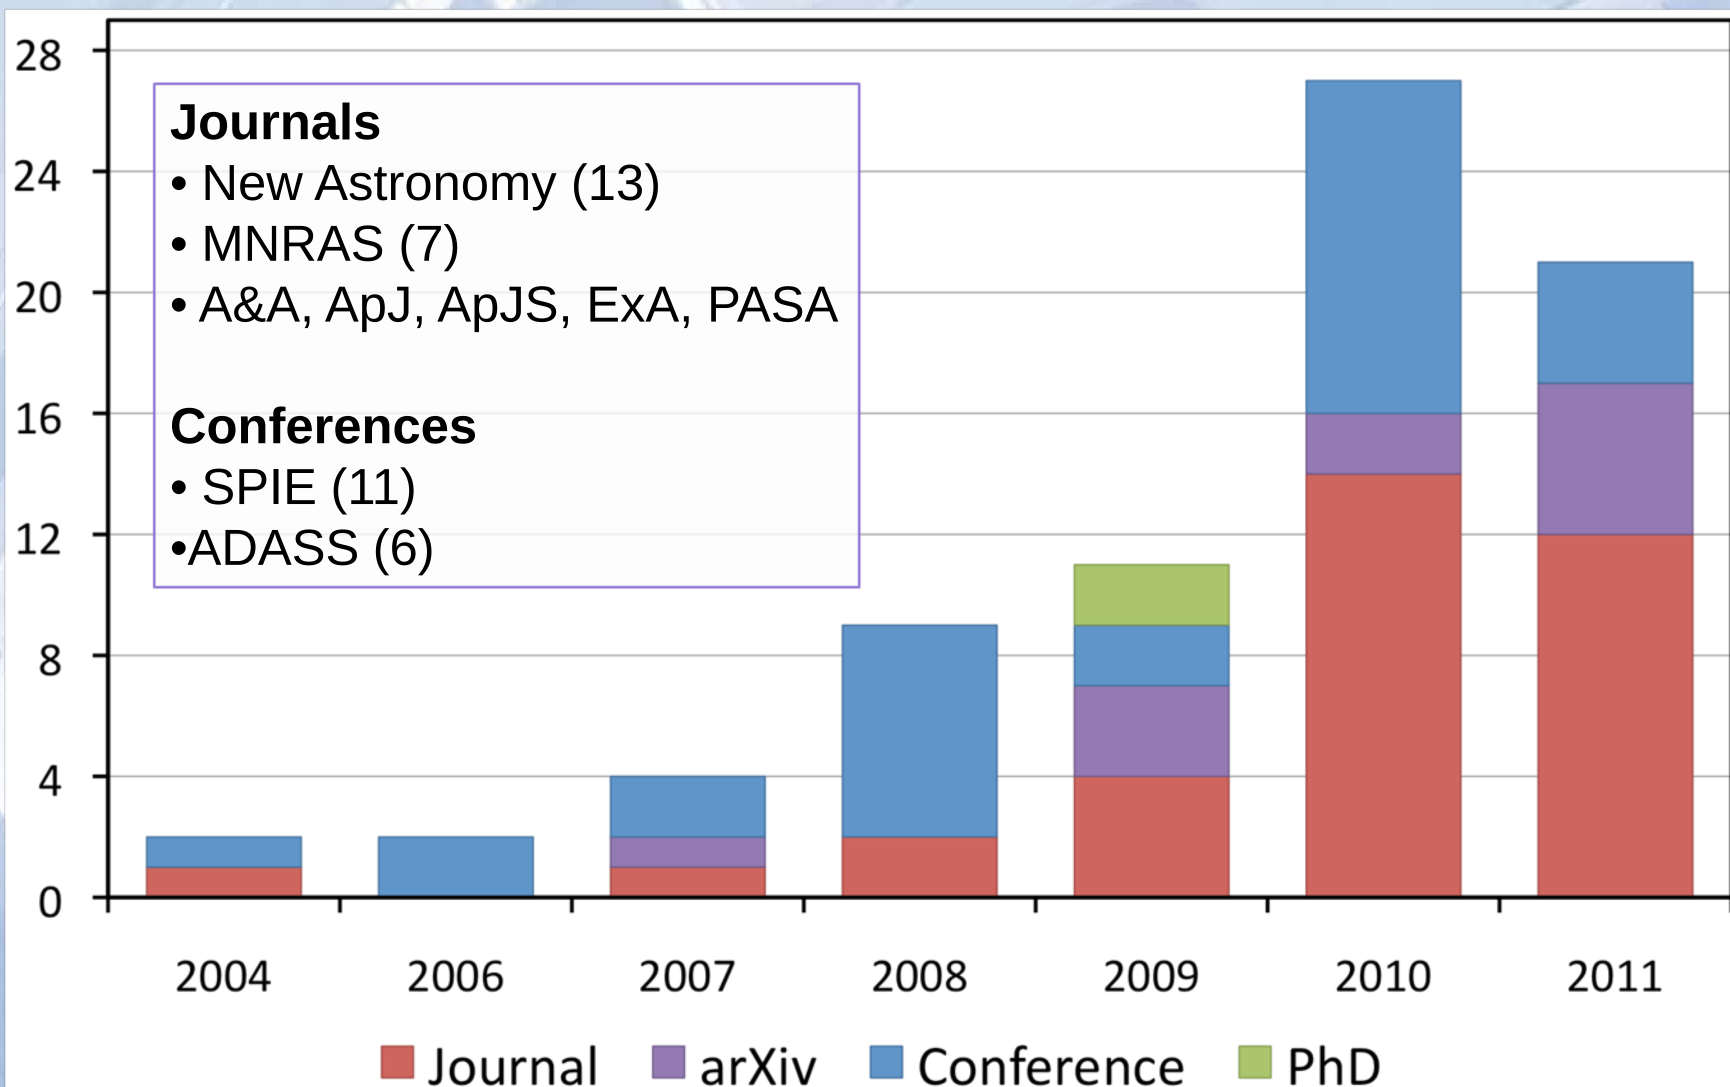

Journals
- New Astronomy (13)
- MNRAS (7)
- A&A, ApJ, ApJS, ExA, PASA

Conferences
- SPIE (11)
- ADASS (6)

39 12 41 2

Other Trends

- Which API?
  - Cg (2; none since 2007)
  - Cuda: 26; since 2008
  - OpenCL: 7 since 2010
- Which card?
  - NVIDIA: 17
    - S1070, C1060, and C2050 cards in six abstracts since 2010
  - ATI: 2
    - Elsen et al. (2007); Pang et al. (2010)

- NVIDIA/CUDA dominance: late appearance of OpenCL?

Reported Speed-ups

- Relative to CPU (mostly single core; a few multi-core)
  - 7x (computing FFT for AO in Rodriguez-Ramos et al. 2006)
  - 600x (solving Kepler's equations in Ford 2009)
  - Most around 10x to 100x or "one-to-two orders of magnitude"
- Caution
  - Why spend time optimising CPU to do a performance test?
  - Single precision vs double precision speed-up?
  - Opportunities to use OpenMP on multicore
- However…GPUs continue to get faster cf. single-core CPUs

TOP500 Supercomputing Sites (June 2011)

Rank	Site	Computer	
1	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	
2	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C NUDT	**GPU**
3	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.	
4	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU Dawning	**GPU**
5	GSIC Center, Tokyo Institute of Technology Japan	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows NEC/HP	**GPU**
6	DOE/NNSA/LANL/SNL United States	Cielo - Cray XE6 8-core 2.4 GHz Cray Inc.	
7	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon 5570/5670 2.93 Ghz, Infiniband SGI	
8	DOE/SC/LBNL/NERSC United States	Hopper - Cray XE6 12-core 2.1 GHz Cray Inc.	
9	Commissariat a l'Energie Atomique (CEA) France	Tera-100 - Bull bullx super-node S6010/S6030 Bull SA	
10	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband IBM	

Source: www.top500.org

The Green500 (June 2011) – Energy Efficiency

Green500 Rank	MFLOPS/W	Site*		Computer*	Total Power (kW)
1	2097.19	IBM Thomas J. Watson Research Center		NNSA/SC Blue Gene/Q Prototype 2	40.95
2	1684.20	IBM Thomas J. Watson Research Center		NNSA/SC Blue Gene/Q Prototype 1	38.80
3	1375.88	Nagasaki University	GPU	DEGIMA Cluster, Intel i5, ATI Radeon GPU, Infiniband QDR	34.24
4	958.35	GSIC Center, Tokyo Institute of Technology	GPU	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows	1243.80
5	891.88	CINECA / SCS - SuperComputing Solution	GPU	iDataPlex DX360M3, Xeon 2.4, nVidia GPU, Infiniband	160.00
6	824.56	RIKEN Advanced Institute for Computational Science (AICS)		K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	9898.56
7	773.38	Forschungszentrum Juelich (FZJ)		QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
8	773.38	Universitaet Regensburg		QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
9	773.38	Universitaet Wuppertal		QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
10	718.13	Universitaet Frankfurt	GPU	Supermicro Cluster, QC Opteron 2.1 GHz, ATI Radeon GPU, Infiniband	416.78

Source: www.green500.org

High Performance Computing with GPU Clusters

- University of Heidelberg
  - Kolob cluster (40 x Tesla C870)
- National Astronomical Observatories of China
  - Silk Road project (170 GPUs)
- Nagasaki University
  - Hamada & Nitadori (2010)
  - 576 x NVIDIA GT200
  - 3 billion particle N-body system
  - 190 Gflop/s for $400,000 USD

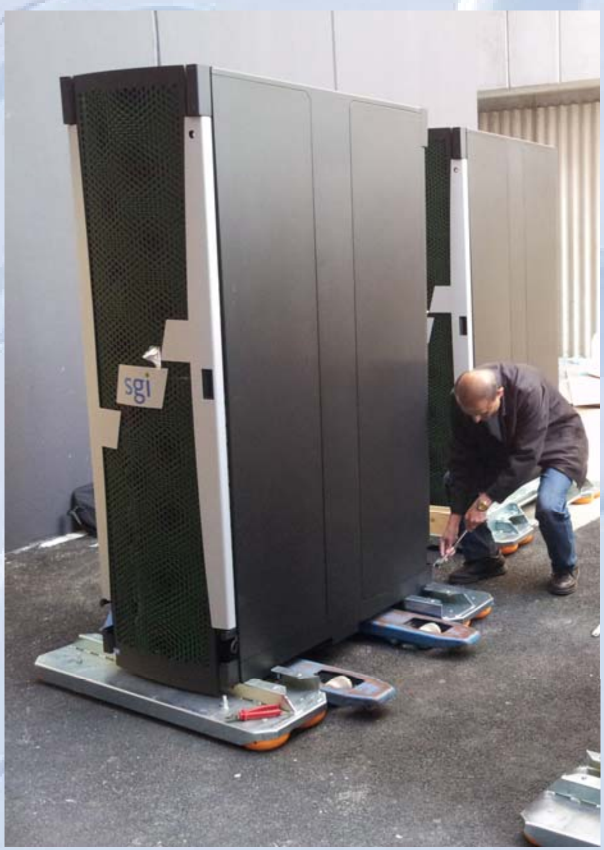

Credit: Gin Tan

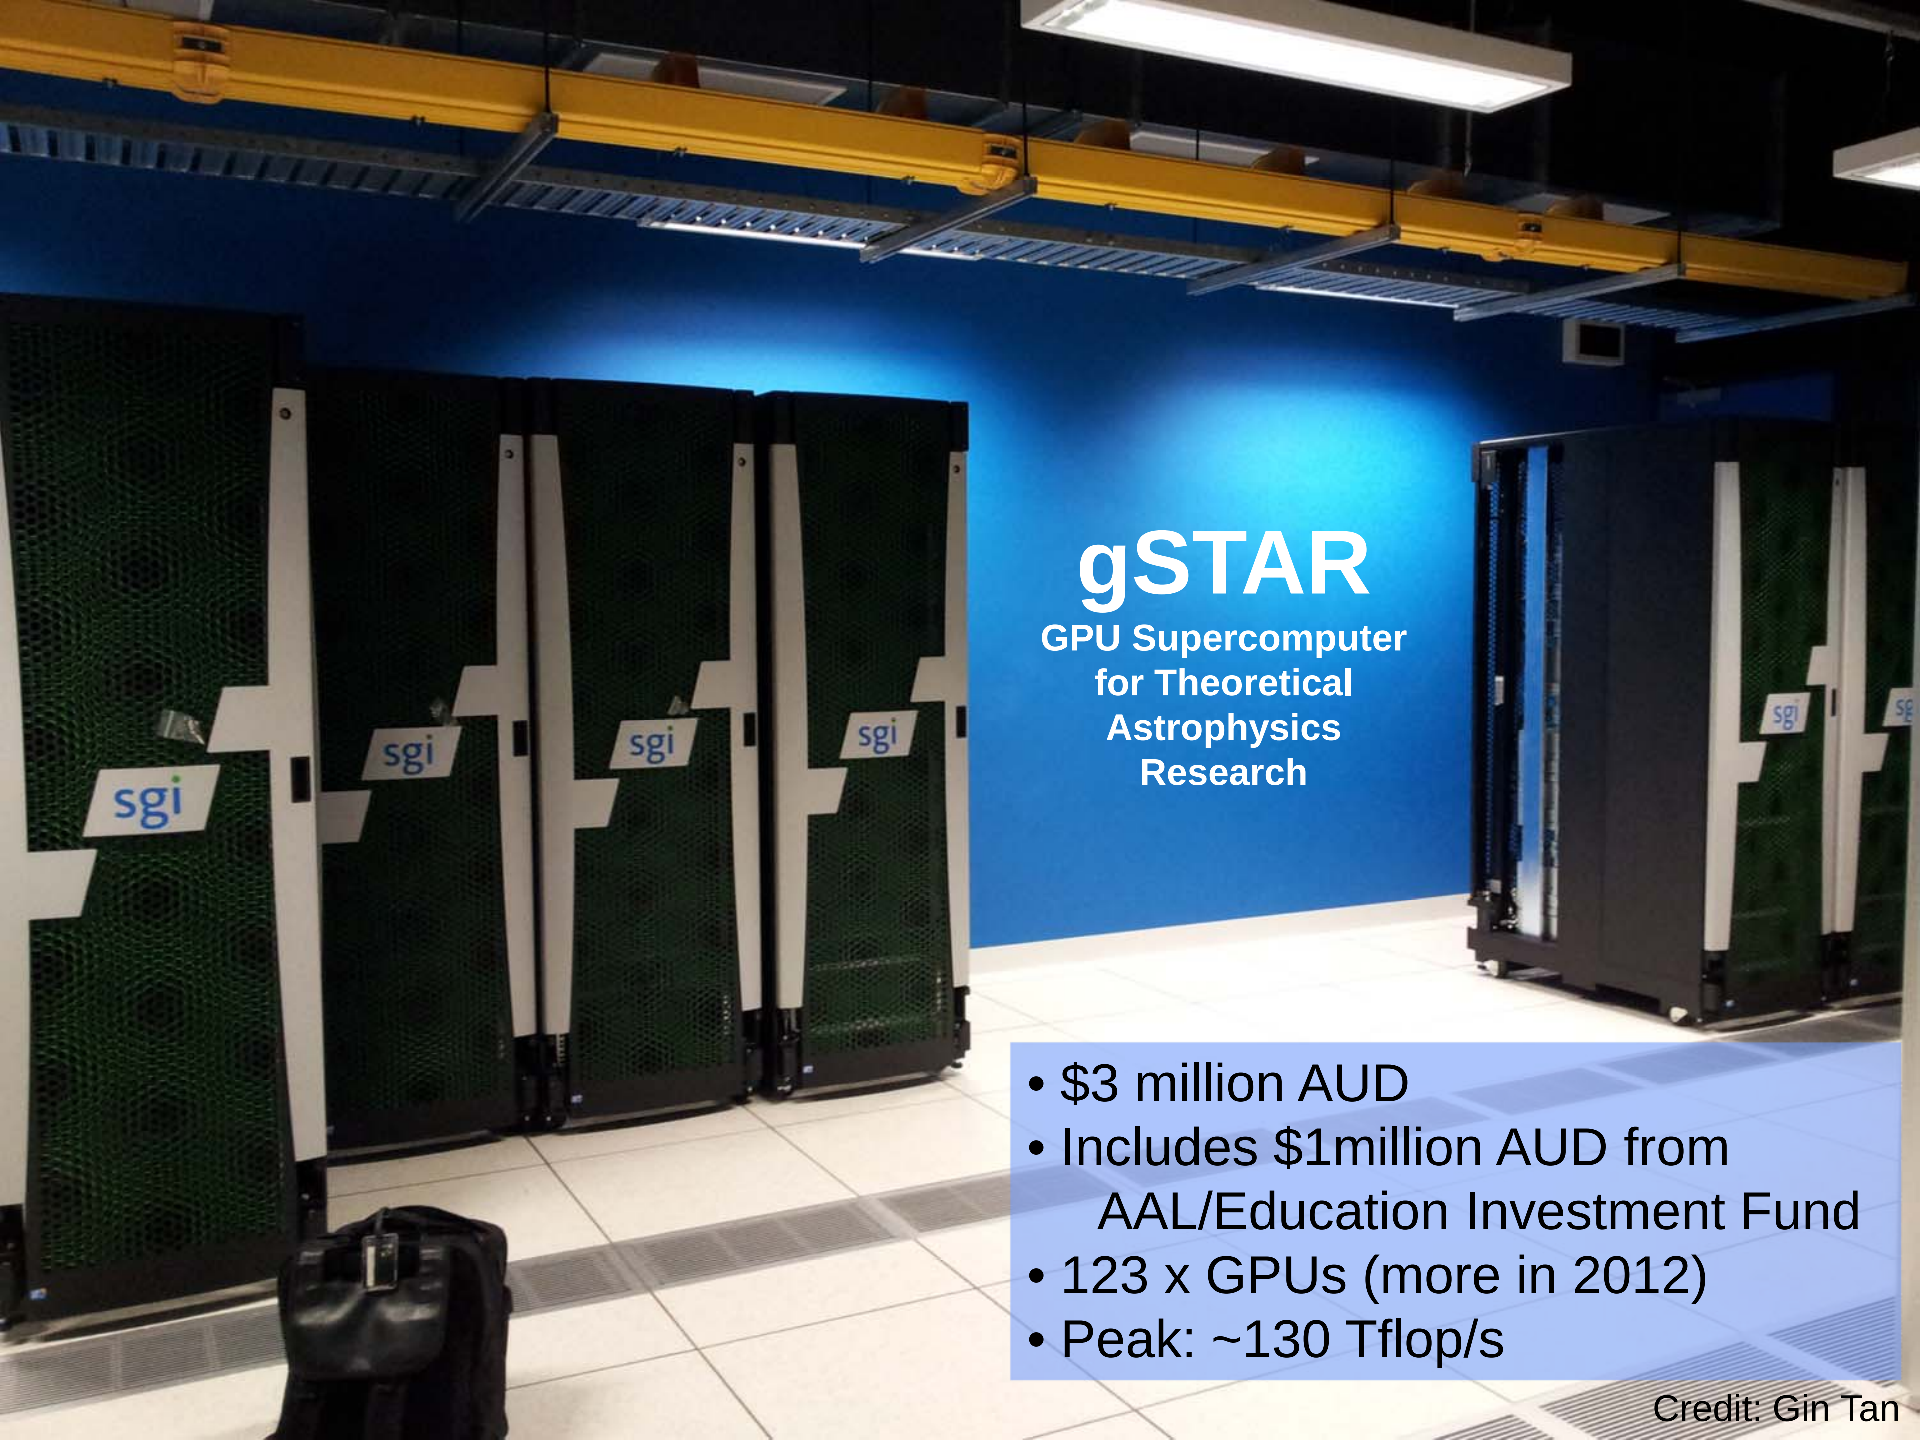
gSTAR
GPU Supercomputer for Theoretical Astrophysics Research
• $3 million AUD
• Includes $1million AUD from AAL/Education Investment Fund
• 123 x GPUs (more in 2012)
• Peak: ~130 Tflop/s
Credit: Gin Tan

Early science on gSTAR

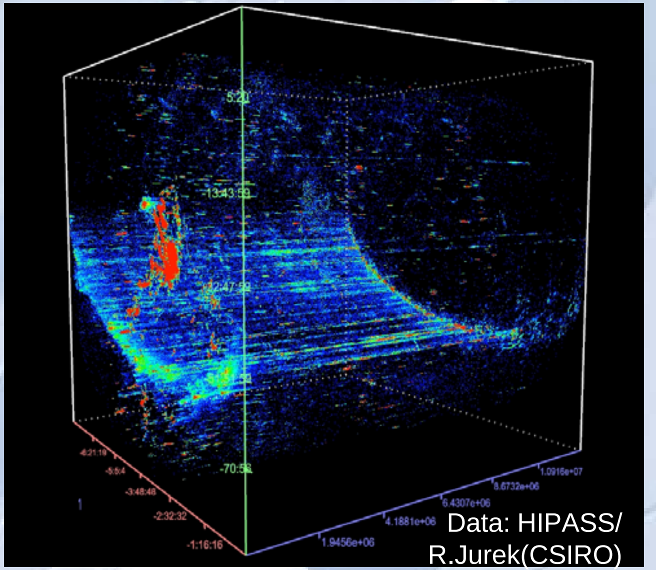

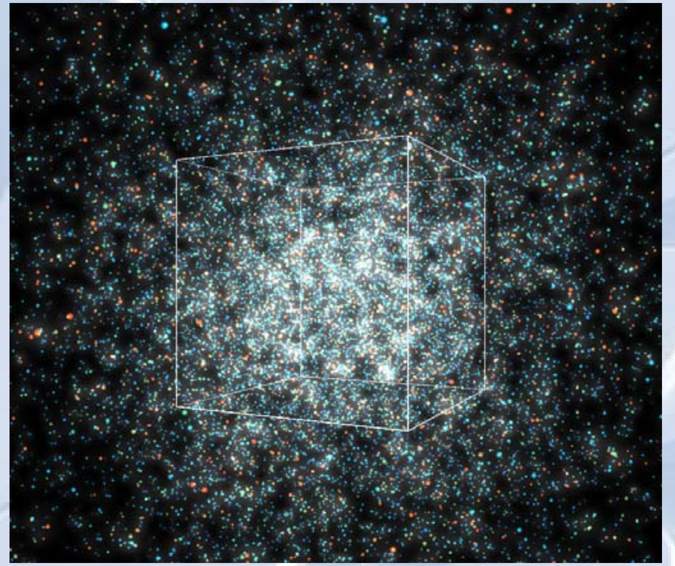

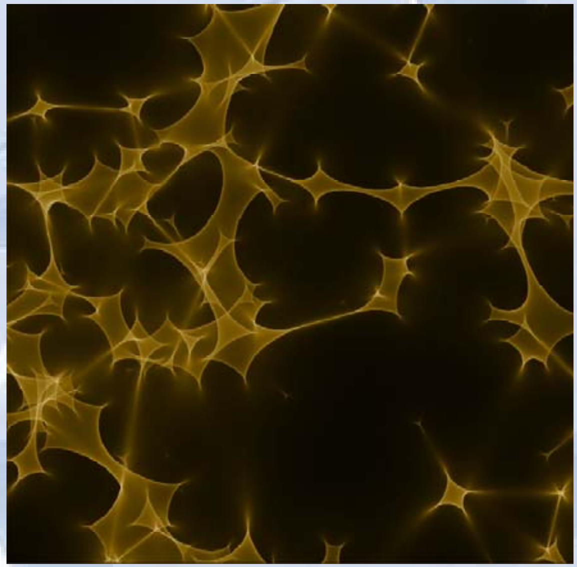

- Real-time, 3D volume rendering of terascale spectral cubes
  - **Hassan,** Fluke, Barnes (Monash)
- Direct N-body star cluster simulations
  - Hurley, **Sippel, Madrid, Moyano-Loyola**
- Gravitational microlensing parameter survey
  - **Vernardos**, Fluke, Bate (Sydney)

Bold = PhD student

Accelerating the Rate of Astronomical Discovery

- Run an individual problem faster
  - Minutes instead of days, weeks instead of months
  - Real-time solutions
    - Wave-front correction
    - Transient detection (Next two talks)
- Run more problems in the same wall time
  - Parameter space exploration
    - Black hole inspirals – Herrmann et al. (2010)
    - Solving Kepler's equations – Ford (2009)
    - Lyman-α forest simulations – Greig et al. (2011)
  - Important use for GPU Clusters
    - Statistical analysis vs. over-analysis?

Accelerating the Rate of Astronomical Discovery

- Solve a bigger problem size in same wall time as smaller problem on CPU
  - Work at higher resolution, more time-steps, etc.
  - Terascale (petascale?) image processing/analysis
  - Data mining
  - However:
    - Does the problem fit in memory? [A.Hassan talk]
    - Bottleneck moves to data transfer

Accelerating the Rate of Astronomical Discovery

- Solve a more complex problem in the same wall time as simpler problem on CPU
  - More accurate solution methods
  - Algorithms with improved accuracy
- Provide much lower price/performance compared to CPU
  - More astronomers able to access Tflop/s HPC

Why aren't we all using GPUs already?

Challenges:

- Cannot run existing code – it must be modified in some way
- Need to identify, implement and optimise *relevant* algorithms
- Parallel programming concepts not as familiar amongst astronomer-programmers
- Can get simple speed-ups on multi-core e.g. OpenMP

Concluding Remarks

- Dawn of the petascale data era
  - New challenges in data processing/simulation
- GPU-powered HPC clusters offer low-cost opportunity to explore new, scalable, massively parallel algorithms
- GPU speed-ups can accelerate the rate of discovery
- The future of computing is here, and it is massively parallel

Here it is again … in parallel

I'll take all of your questions simultaneously…

ACCELERATING THE RATE OF ASTRONOMICAL DISCOVERY WITH GPU-ENABLED CLUSTERS

SWINBURNE UNIVERSITY OF TECHNOLOGY

NVIDIA CUDA RESEARCH CENTER

Dr Christopher Fluke

Scientific Computing & Visualisation Group

ADASS 2011

Thanks to B.Barsdell (Swin), A.Hassan (Swin), D.Barnes (Monash) and ADASS POC

Bonus Slides

gSTAR: Specification

- 51 dual-socket compute nodes each with 2 GPUs
  - NVIDIA C2070: 6GB RAM
- 3 high-density nodes each with 7 GPUs
  - M2090: 6GB RAM
- >1.0 PB disk space (Lustre file system)
- QDR InfinbandB (non-blocking)
- ~130 Tflop/s (theoretical peak)
- Phase 2: more GPUs next year

Credit: Gin Tan

Methods (82/94):

- Demonstrate that an algorithm is suited to GPU
- Quote a speed-up or peak processing performance

Applications (9/94):

- Use a GPU code to achieve new science result

Philosophy (3/94):

- Adoption of GPUs for scientific computing in astronomy

Top500 Supercomputing Sites (June 2011)

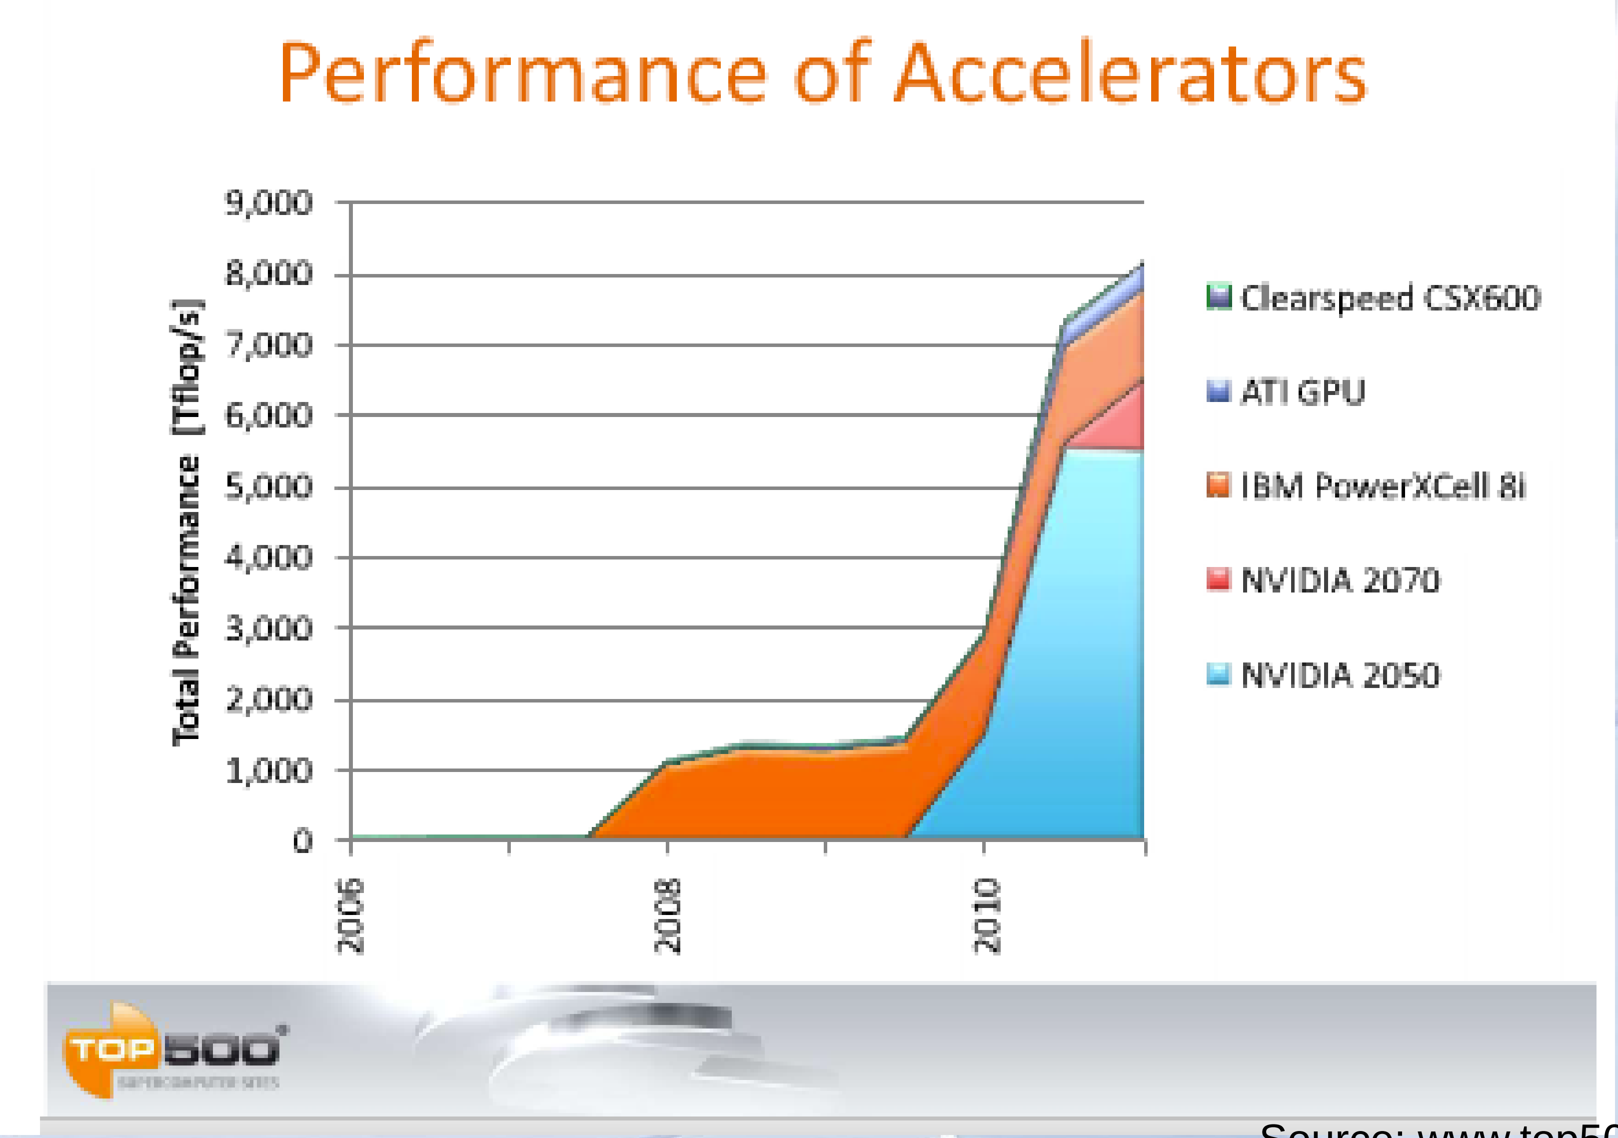

Source: www.top500.org

Top500 Supercomputing Sites (June 2011)

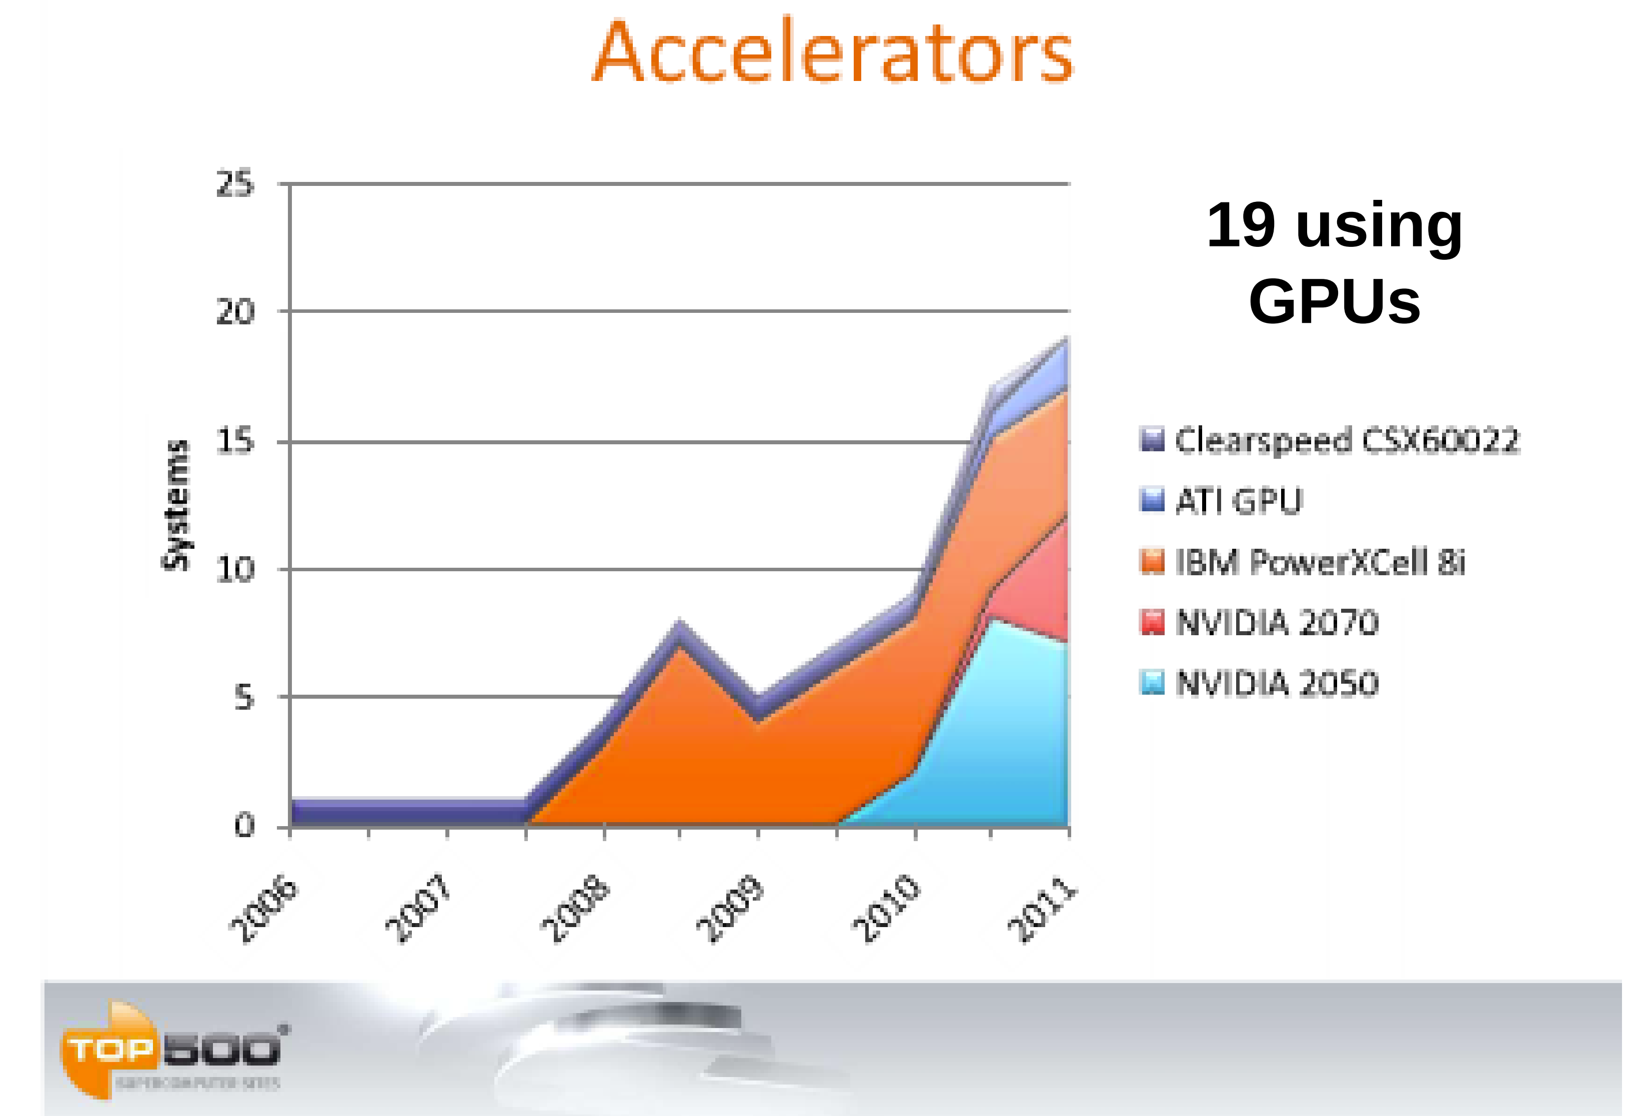

19 using GPUs

Source: www.top500.org

GPUs @ Swinburne

- Adoption and Applications: Ben Barsdell, David Barnes
- Visualisation: Amr Hassan
- Gravitational Lensing: Giorgos Vernardos, Nick Bate, Alex Thompson

---

- Pulsars: Matthew Bailes, Jonathon Kocz, Paul Coster, Willem van Straten, Ben Barsdell
- Cosmology: Darren Croton, Max Berynk
- N-body simulations: Juan Madrid, Anna Sippel, Guido Moyano Loyola, Jarrod Hurley

---

Disclaimer:

To date, I have written one OpenCL kernel myself. It slowed my code down by a factor of 5. There is nothing wrong with getting other people to write GPU code for you!

Analysing algorithms for GPUs and beyond

B.Barsdell, D.Barnes (Monash), C.Fluke

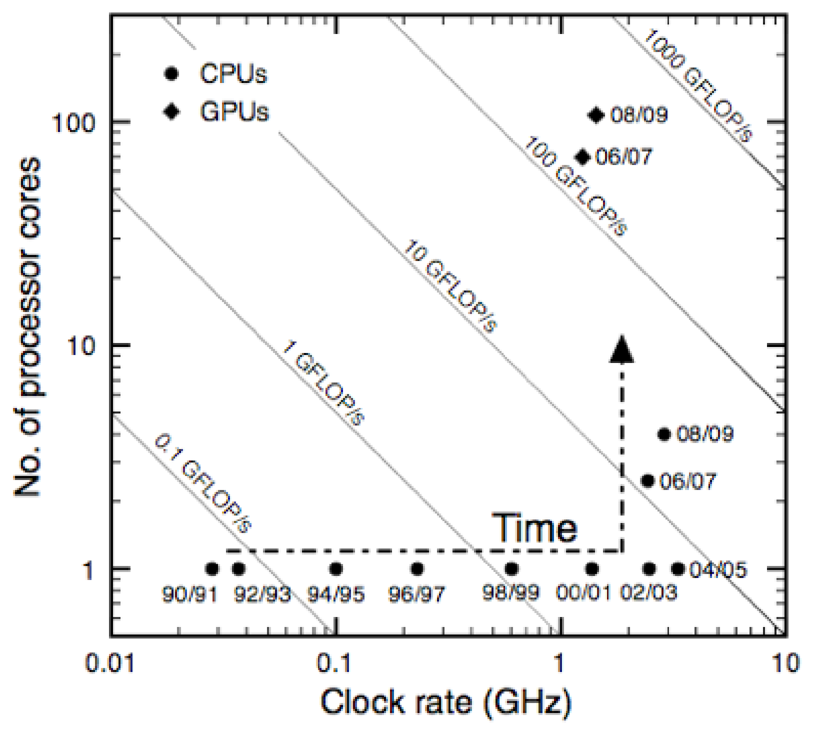

GPUs are taking us to exciting new territories, beyond the current CPU multi-core corner

- Aim: Develop a generalised approach to using GPUs for scientific computing.
- Method: Algorithm analysis techniques allow rapid assessment of GPU-suitability for a broad range of problems.

- A generalised approach to GPUs makes it easier to exploit their power and avoids the risk of wasted development time.

Flynn's Taxonomy

Single instruction Single data | **Single instruction Multiple data** | **Multiple instruction Multiple data**

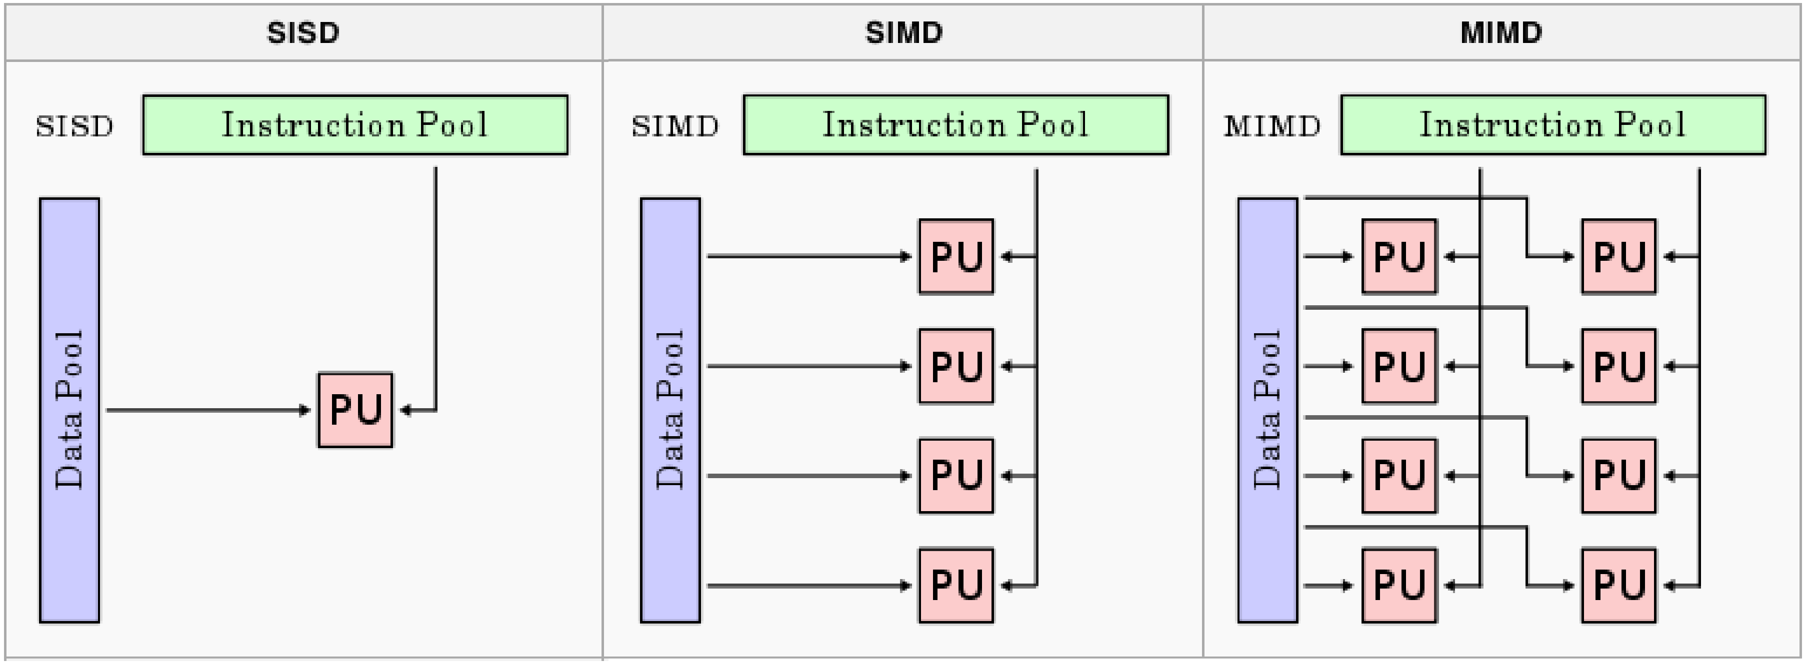

Image: Wikimedia commons

Single core CPU | **GPU** | **Distributed cluster**

Real-time N-Body simulation (+ visualisation)

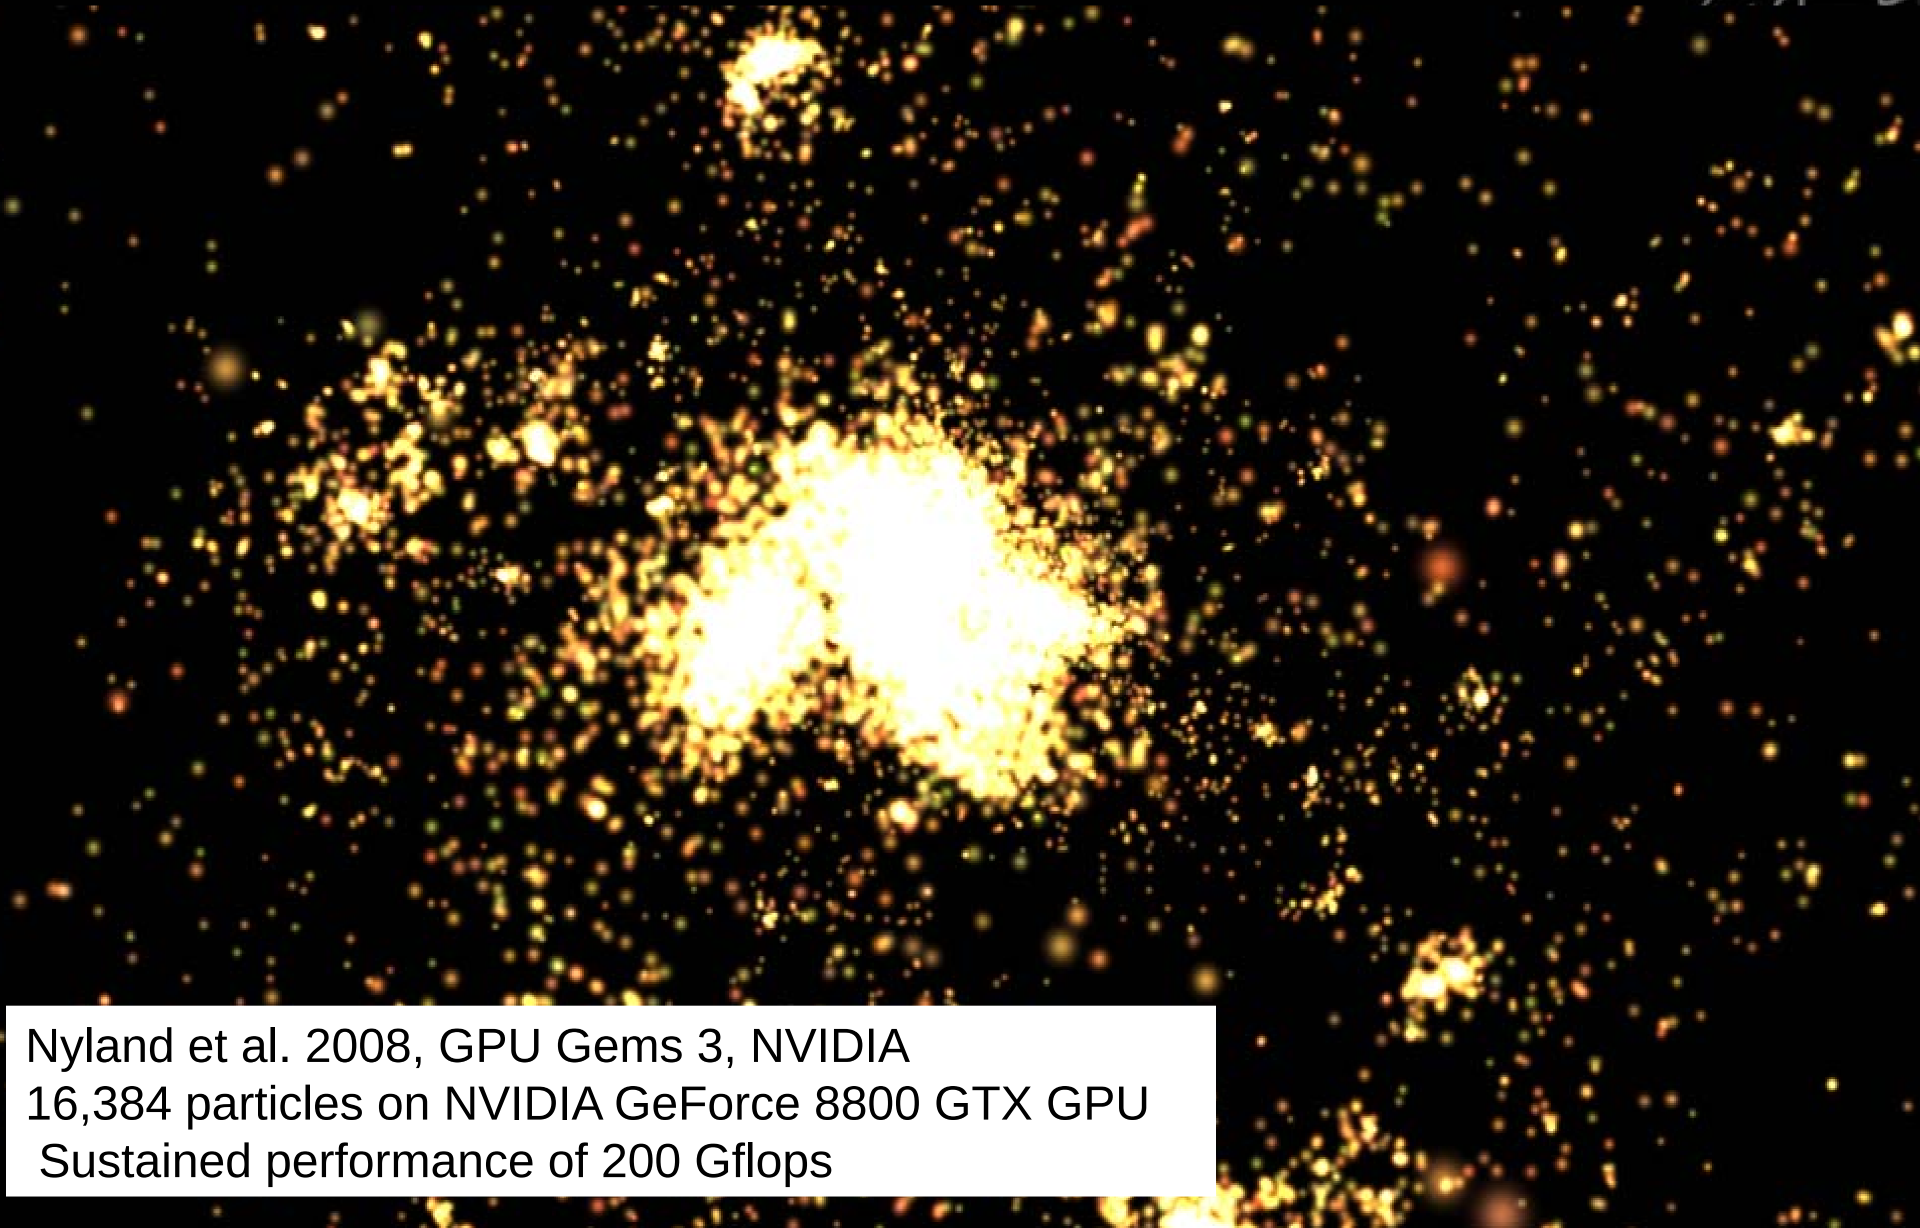

Nyland et al. 2008, GPU Gems 3, NVIDIA
16,384 particles on NVIDIA GeForce 8800 GTX GPU
Sustained performance of 200 Gflops

Records

- Desktop:
  - 1.28 TFLOP/s
  - 4 GPUs in Tesla S1070 (Thompson et al. 2010)
- Cluster:
  - 190 Tflop/s on GPU cluster (Hamada & Nitadori 2010)
- Caution:
  - How to count FLOPS accurately?
  - Mismatch between operations and clock-cycles
  - Rare to get theoretical peak
    - Requires dual issue of multiply + add
- High Performance Computing (HPC) with GPU Clusters

Typical GPU Architecture

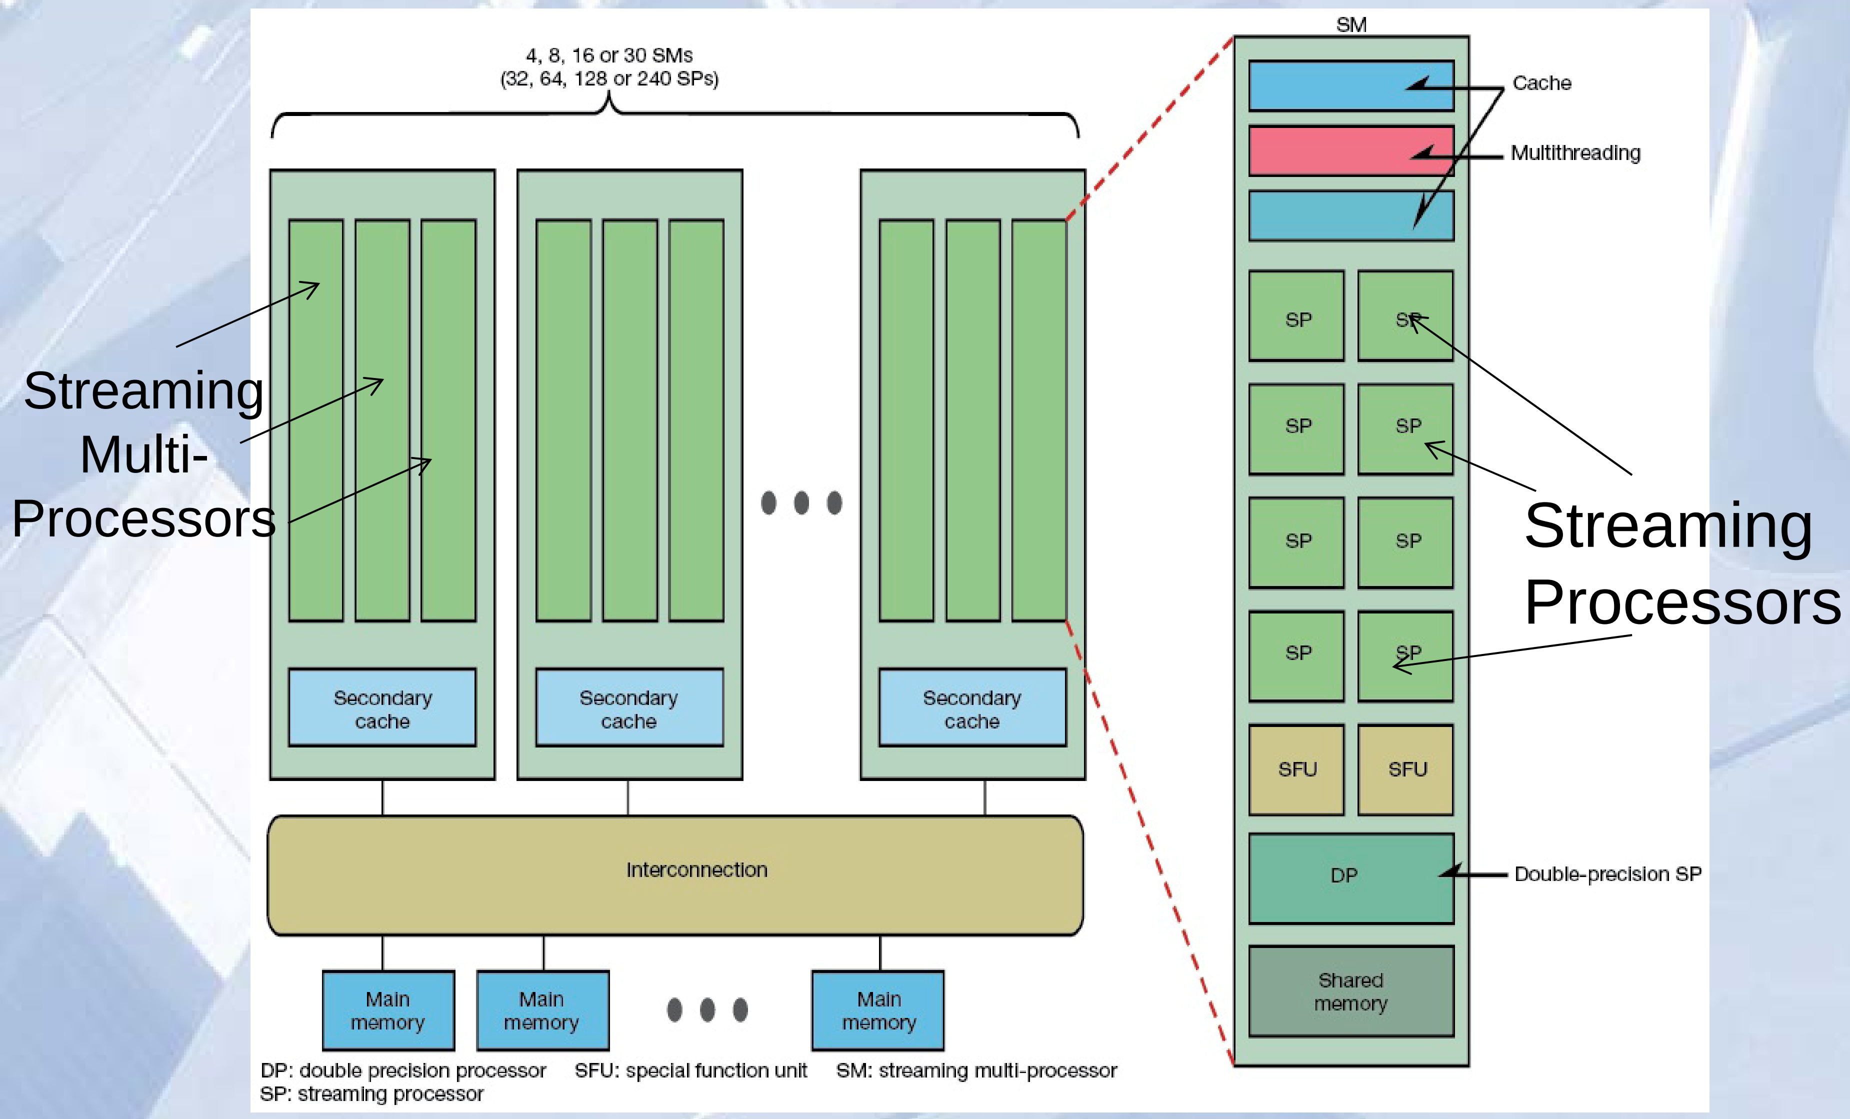

Image: http://techon.nikkeibp.co.jp/article/HONSHI/20090119/164259/