How energy efficiency is driving the future of computing

Simon McIntosh-Smith

University of Bristol, UK

simonm@cs.bris.ac.uk

University of BRISTOL

The real Moore's Law

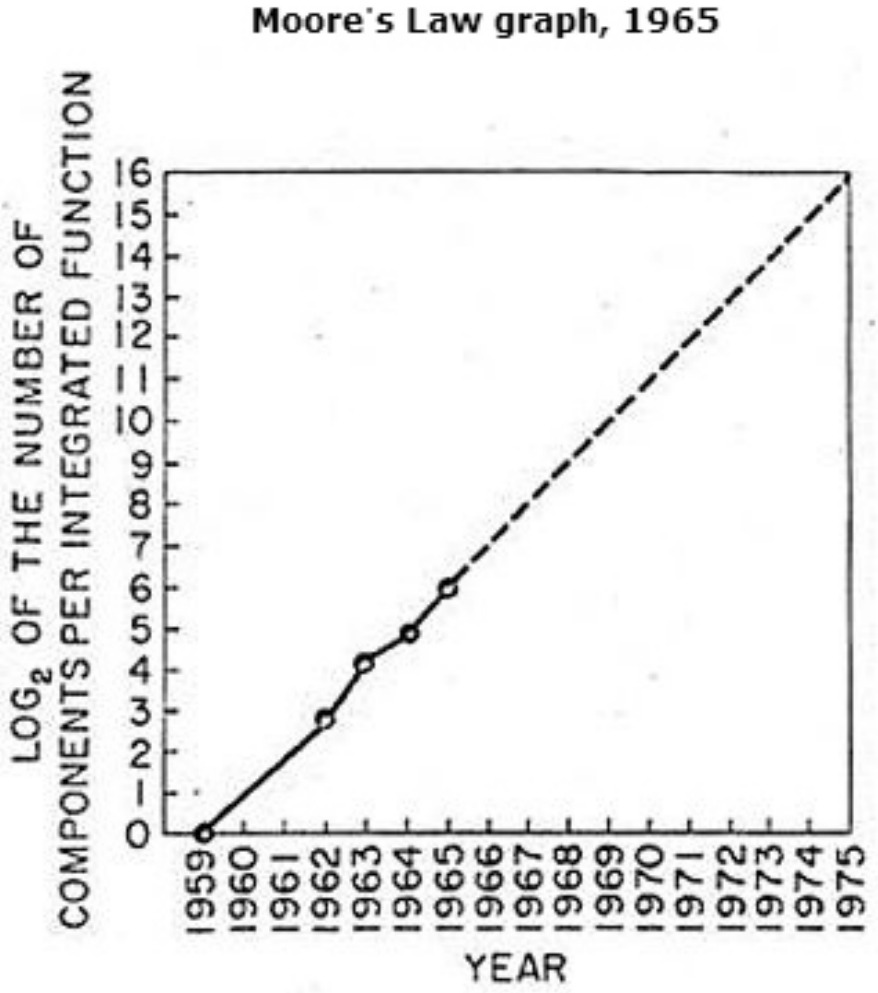

Fig. 2 Number of components per integrated function for minimum cost per component extrapolated vs time.

45 years ago, Gordon Moore observed that the number of transistors on a single chip was doubling rapidly

http://www.intel.com/technology/mooreslaw/

Moore's Law today

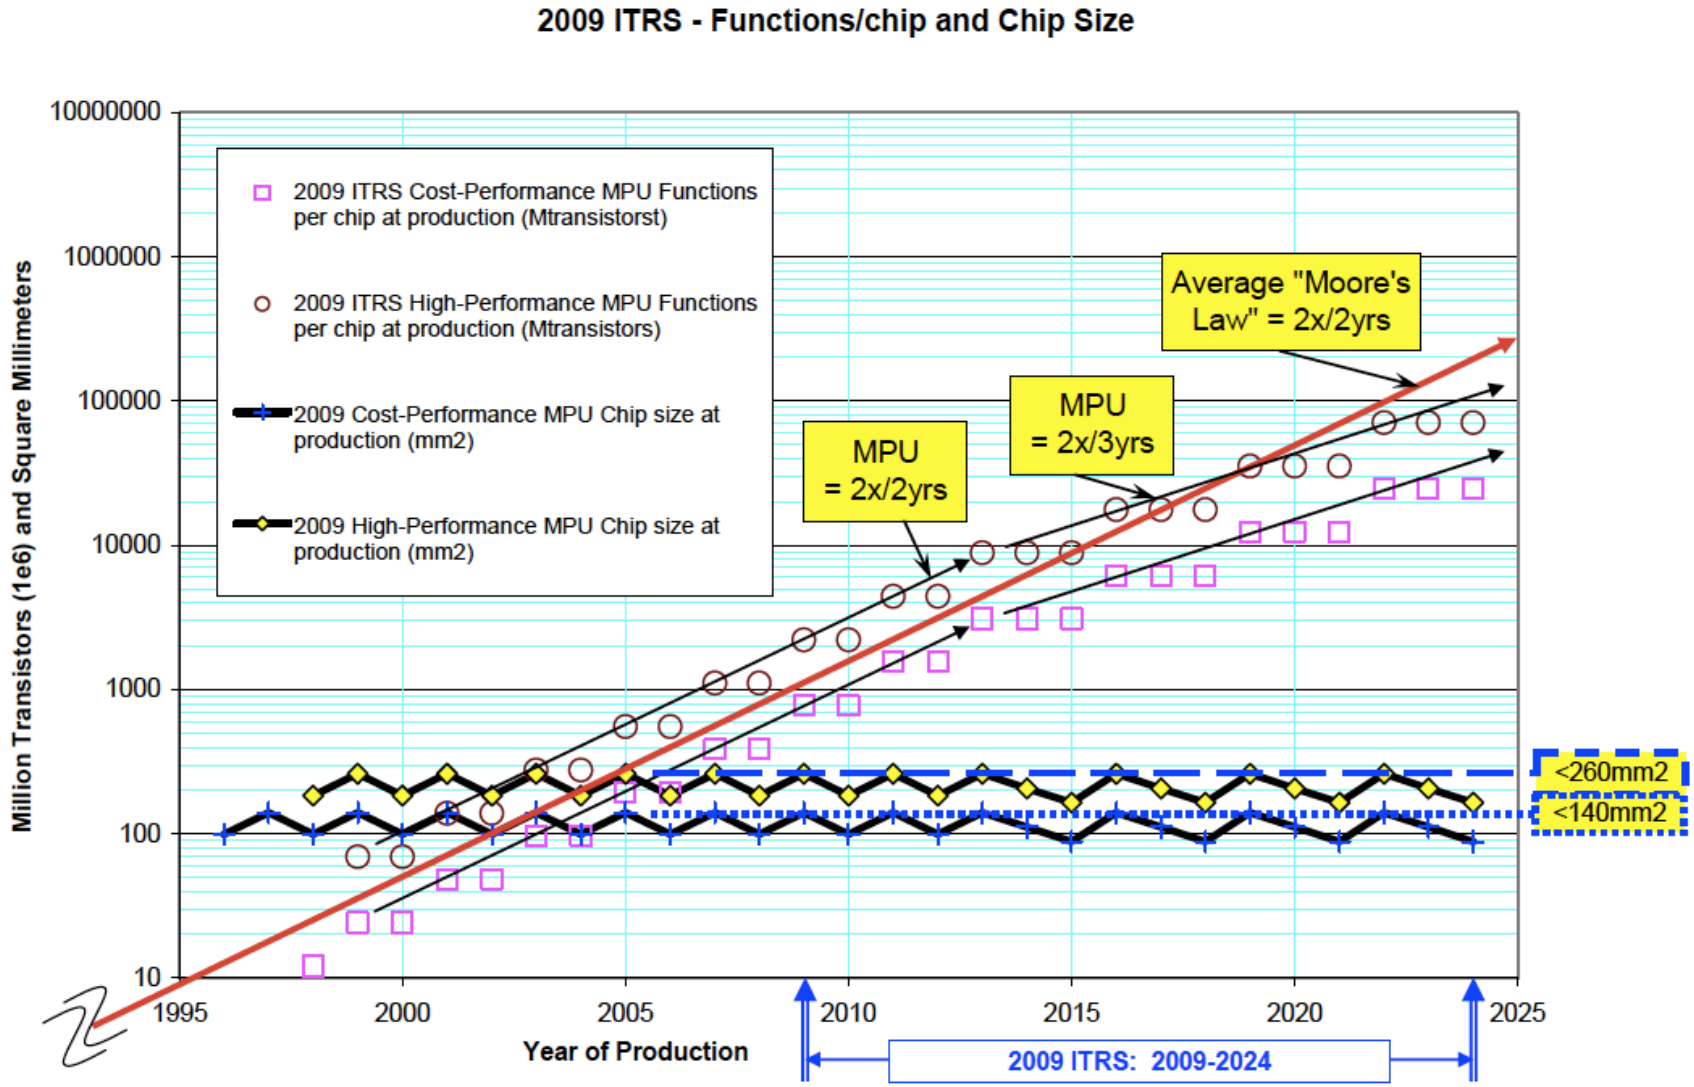

University of BRISTOL

http://www.itrs.net/Links/2009ITRS/2009Chapters_2009Tables/2009_ExecSum.pdf

Moore's Law today

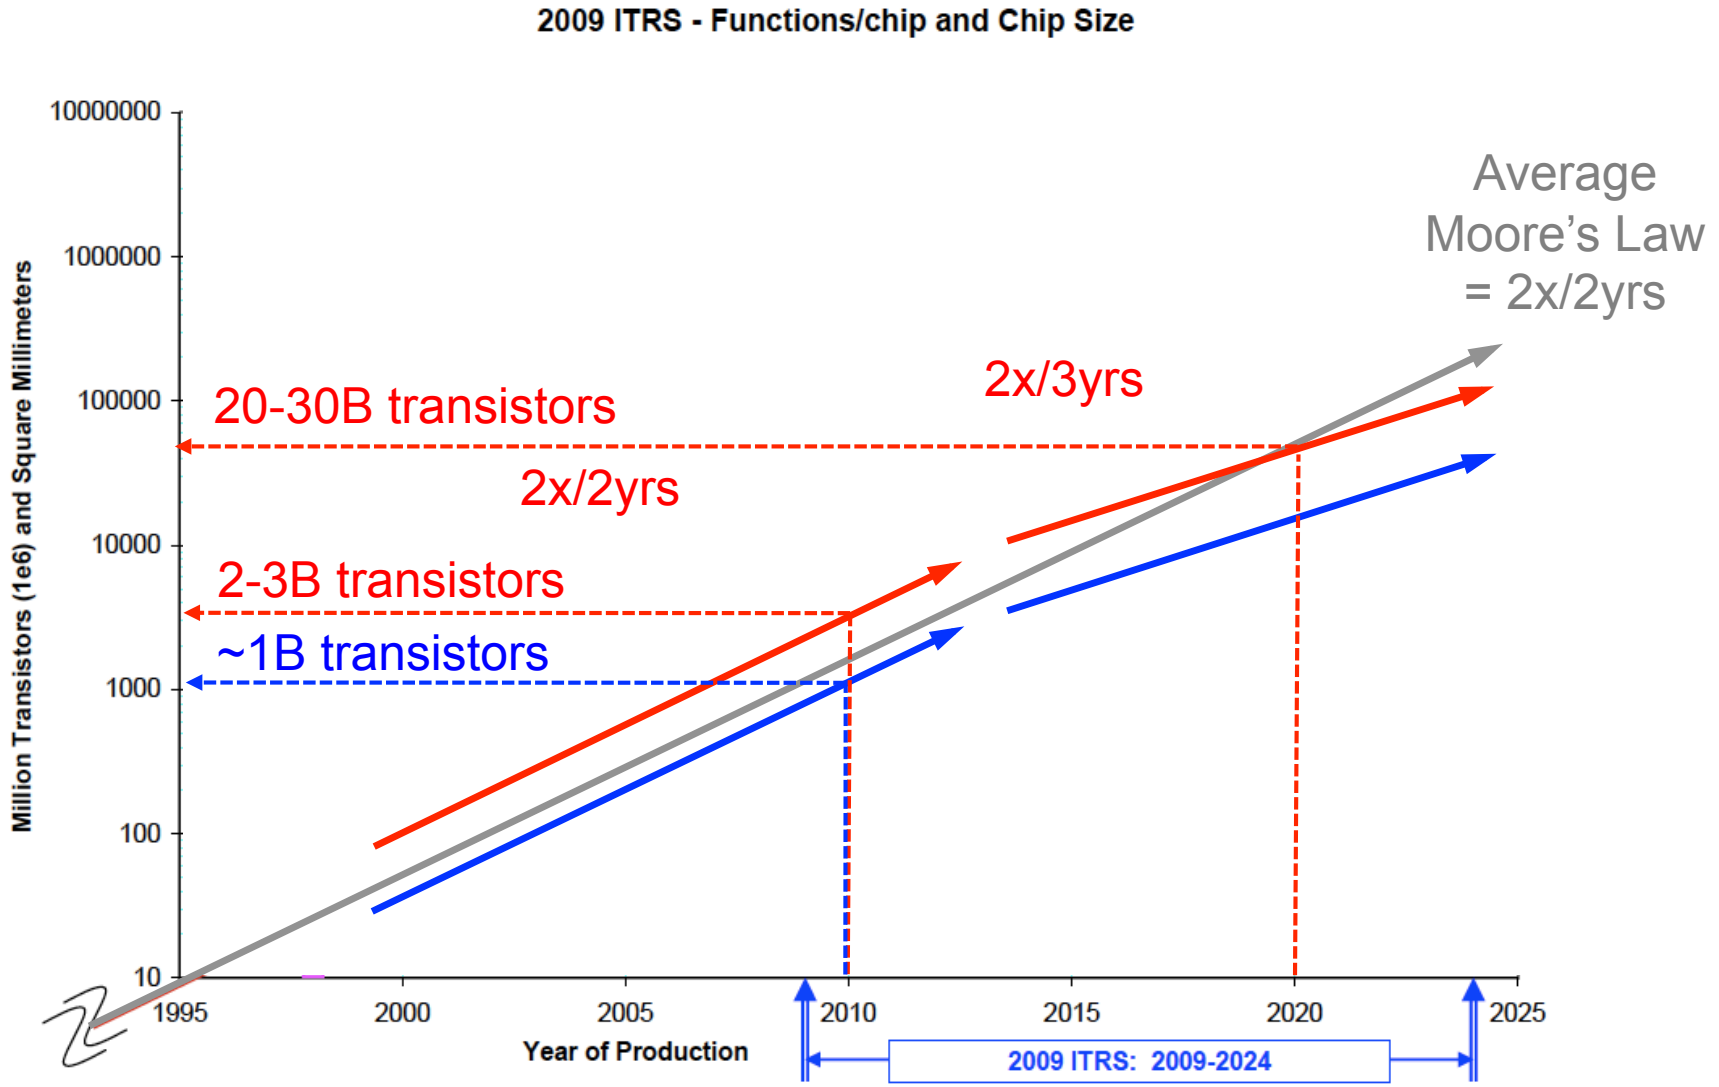

University of BRISTOL

http://www.itrs.net/Links/2009ITRS/2009Chapters_2009Tables/2009_ExecSum.pdf

Important technology trends

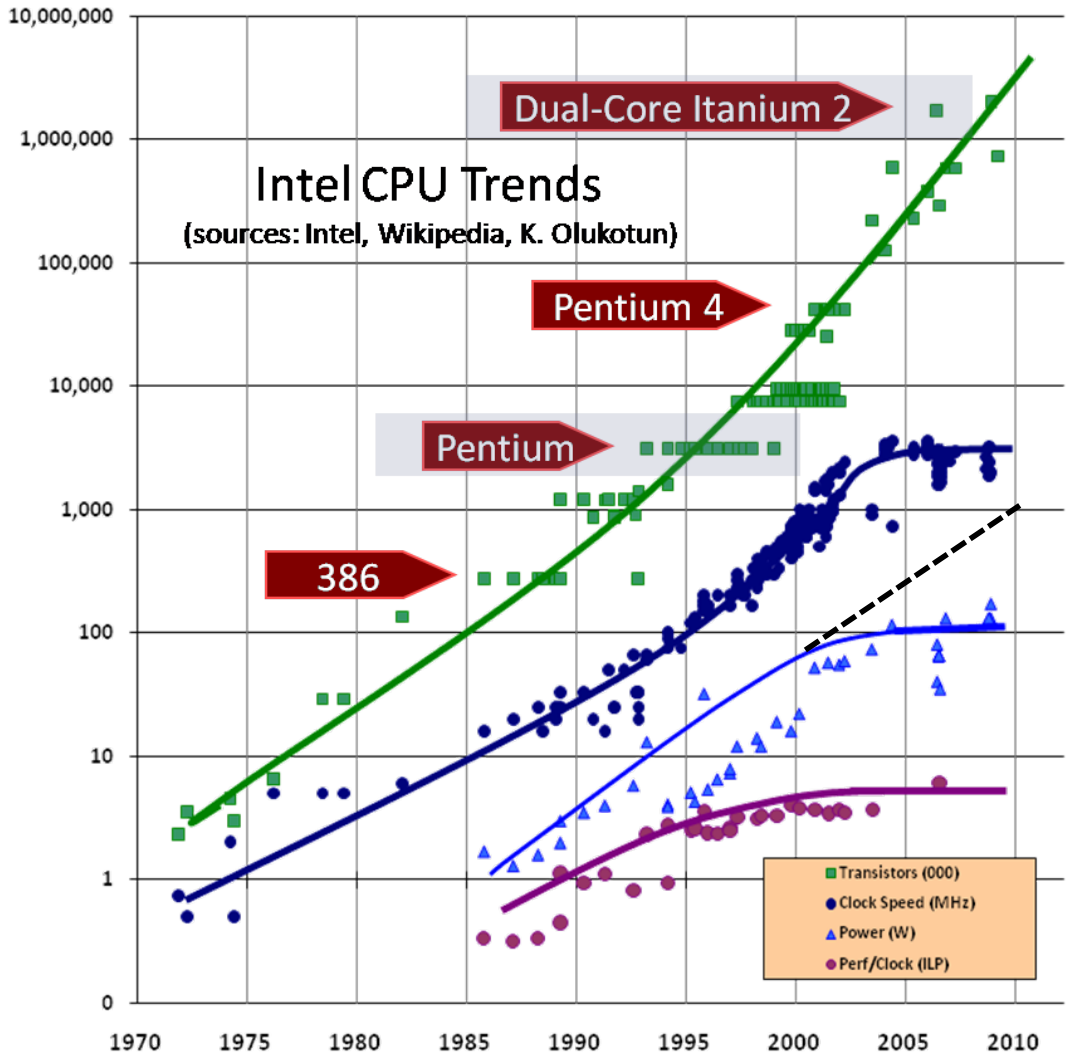

The real Moore's Law

The clock speed plateau

The power ceiling

Instruction level parallelism limit

University of BRISTOL

Herb Sutter, "The free lunch is over", Dr. Dobb's Journal, 30(3), March 2005. On-line version, August 2009.
http://www.gotw.ca/publications/concurrency-ddj.htm

Power-limited regimes

- Processor power consumption now has an upper bound which may even reduce over time

- Power consumption $\propto$
  - Clock frequency
  - Number of transistors (chip area)
    - Number of cores
  - Voltage^2

- When power has an upper bound, "performance per watt = performance"

What to do with billions of transistors?

- Lots more cores on-chip
  - Core designs will stay roughly the same or get simpler
- Power consumption **has to be held in check**
  - Chip voltages ***can't be dialled down any more*** (0.7V)
  - Clock speeds will tend to ***decrease***
  - Memory bandwidth per core will tend to ***decrease***
  - Memory per core will tend to ***decrease***
  - The **DARK SILICON** challenge
- Different types of cores
  - ***Heterogeneous computing***
  - E.g. a few heavyweight (x86) cores together with lots of lightweight (GPU) cores
  - Likely to go the same way as vector units, i.e. mainstream

Future processor architectures

- Hundreds or thousands of cores per chip
- Lower clock speeds
- Highly integrated (mainstream)

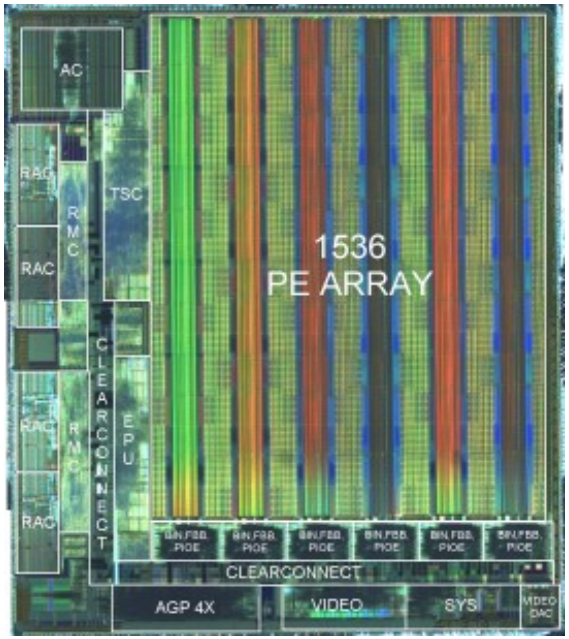

Pixelfusion F150
1,536 simple PEs
200MHz
Circa 2000

University of BRISTOL

The future is now…

AMD's first "Fusion" chip, disclosed at ISSCC in San Francisco, Feb 2010

- 'Llano' Accelerated Processing Unit (APU)
- Integrates a quad core x86 CPU with an OpenCL programmable GPU in the same chip
- Also Intel, Nvidia, IBM…

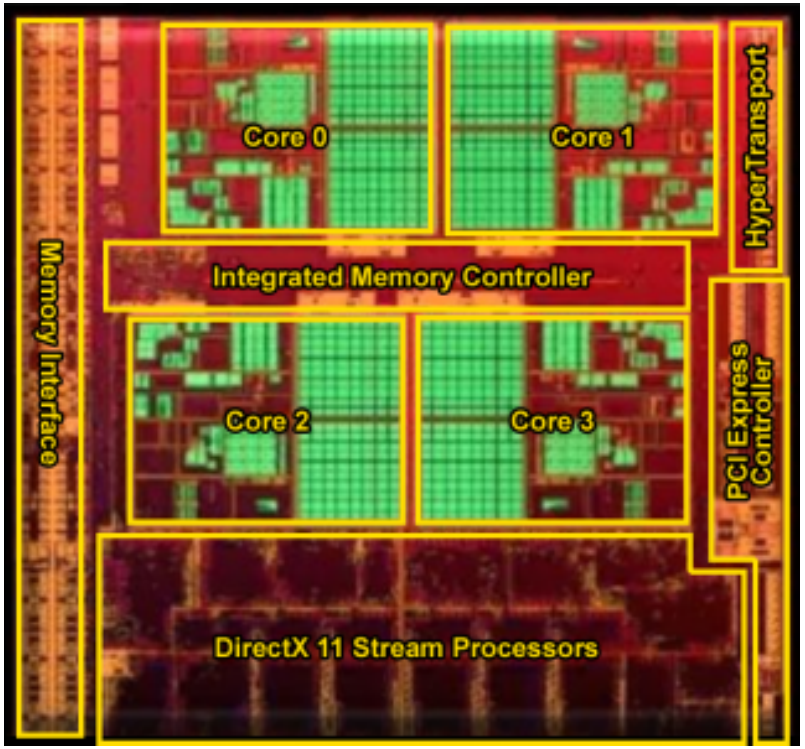

Relative hardware trends

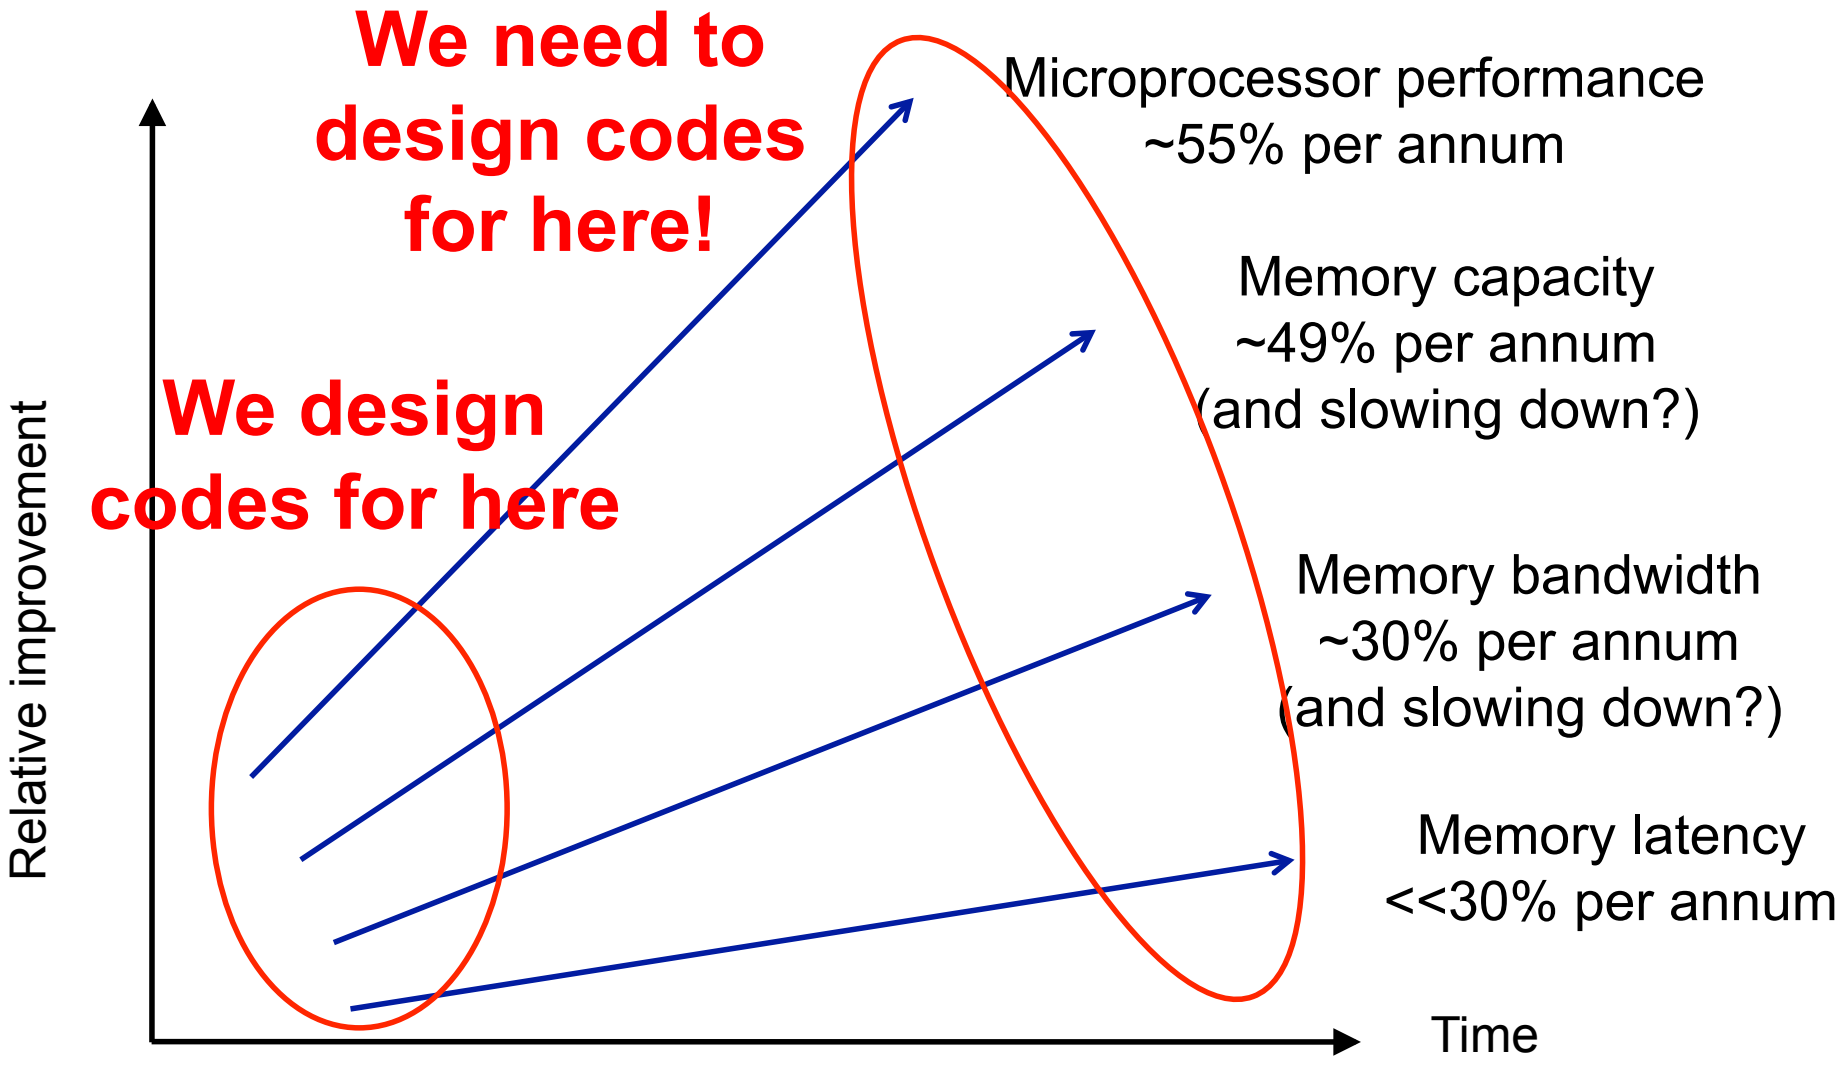

3D stacked memories

- Vertically stack many-core processors with DRAM → *greater bandwidth* and *greater energy efficiency*

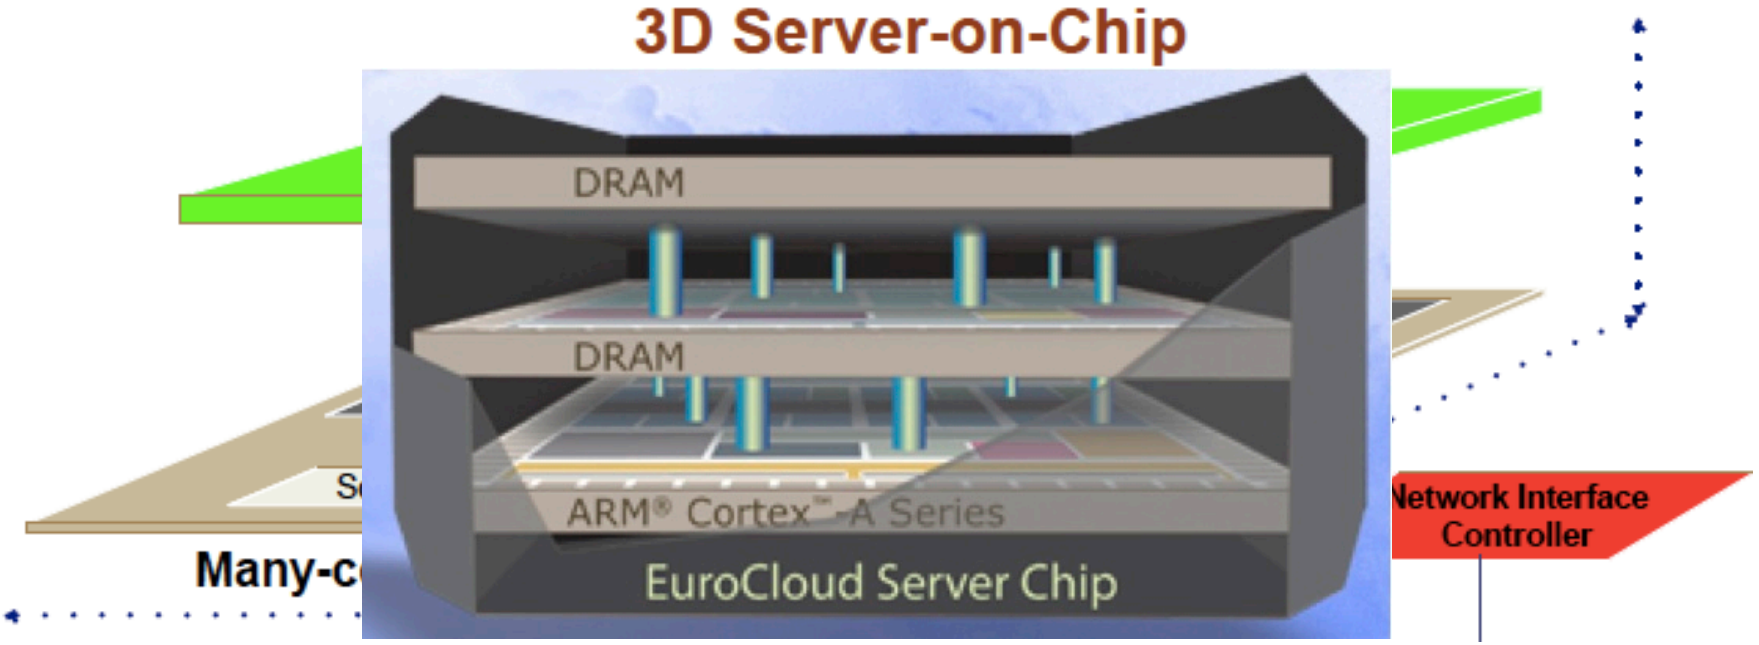

3D stacked memories

- Vertically stack many-core processors with DRAM → *greater bandwidth* and *greater energy efficiency*

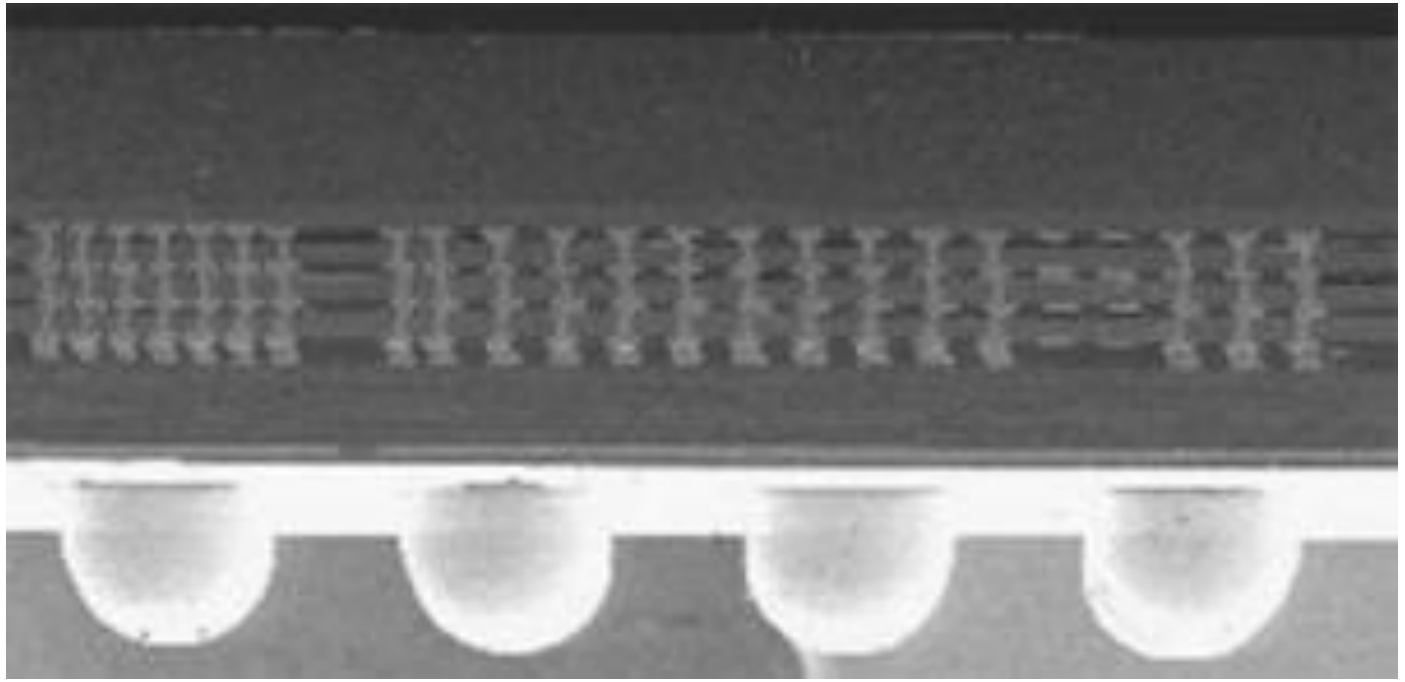

Photonic networks

- Roadmaps to achieve ExaFLOPs (10^{18}) by 2018 are relying on some major hardware breakthroughs to improve energy efficiency
- Prof Keren Bergmen's work at Columbia sponsored by US DoE, Intel, IBM
- ***Moving data*** becoming an increasingly dominant fraction of energy dissipation in microelectronics
  - "Compute free, bandwidth expensive"

Photonic NoC integration

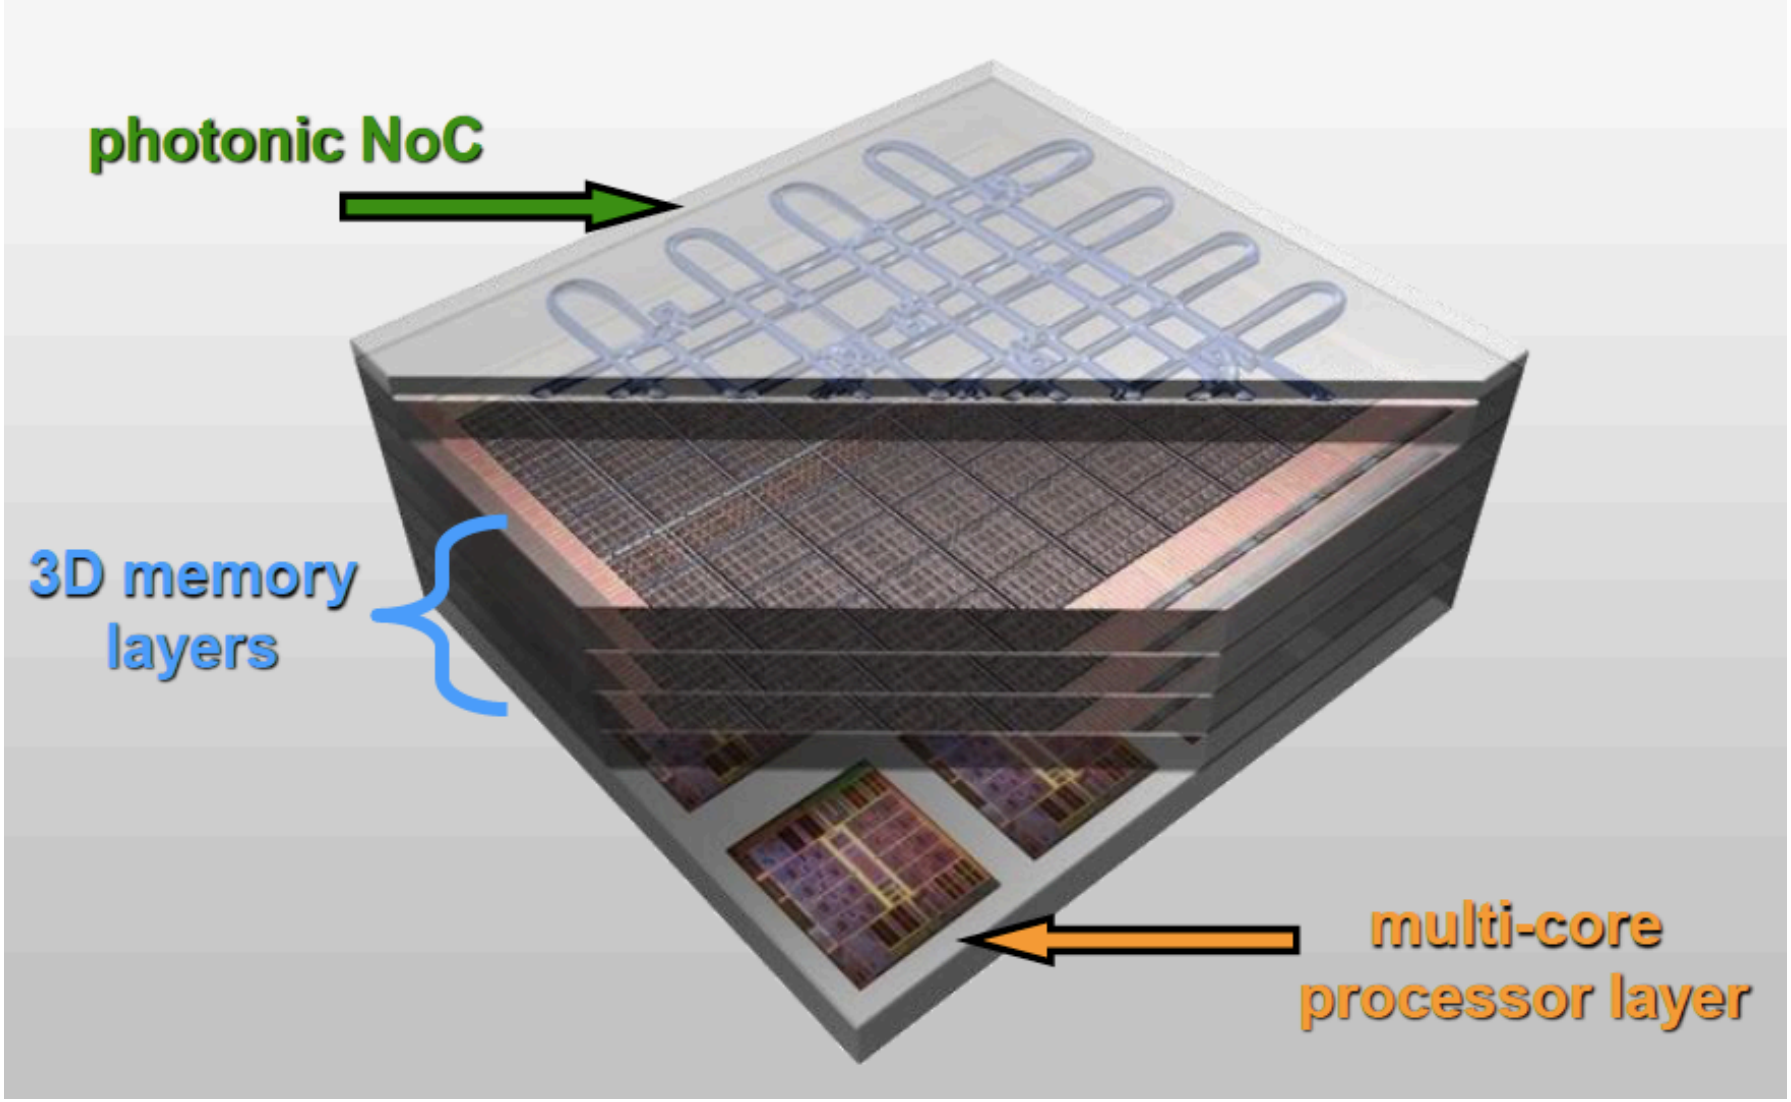

Optically interconnected supercomputing board

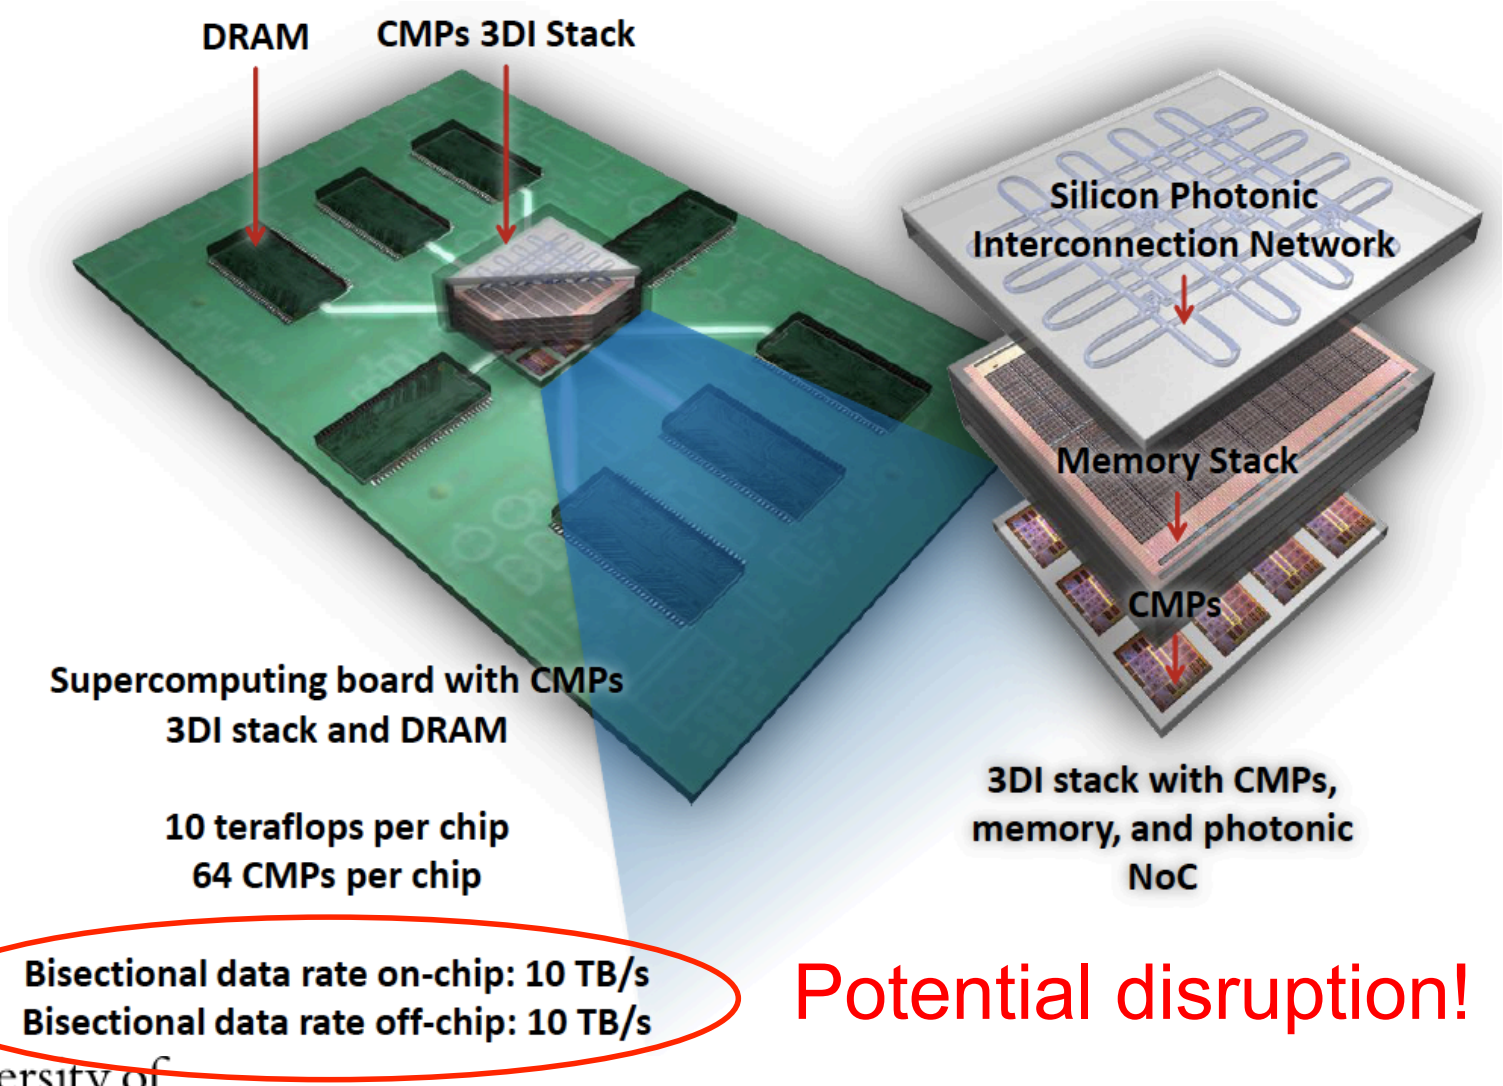

Potential disruption!

Silicon photonic integration

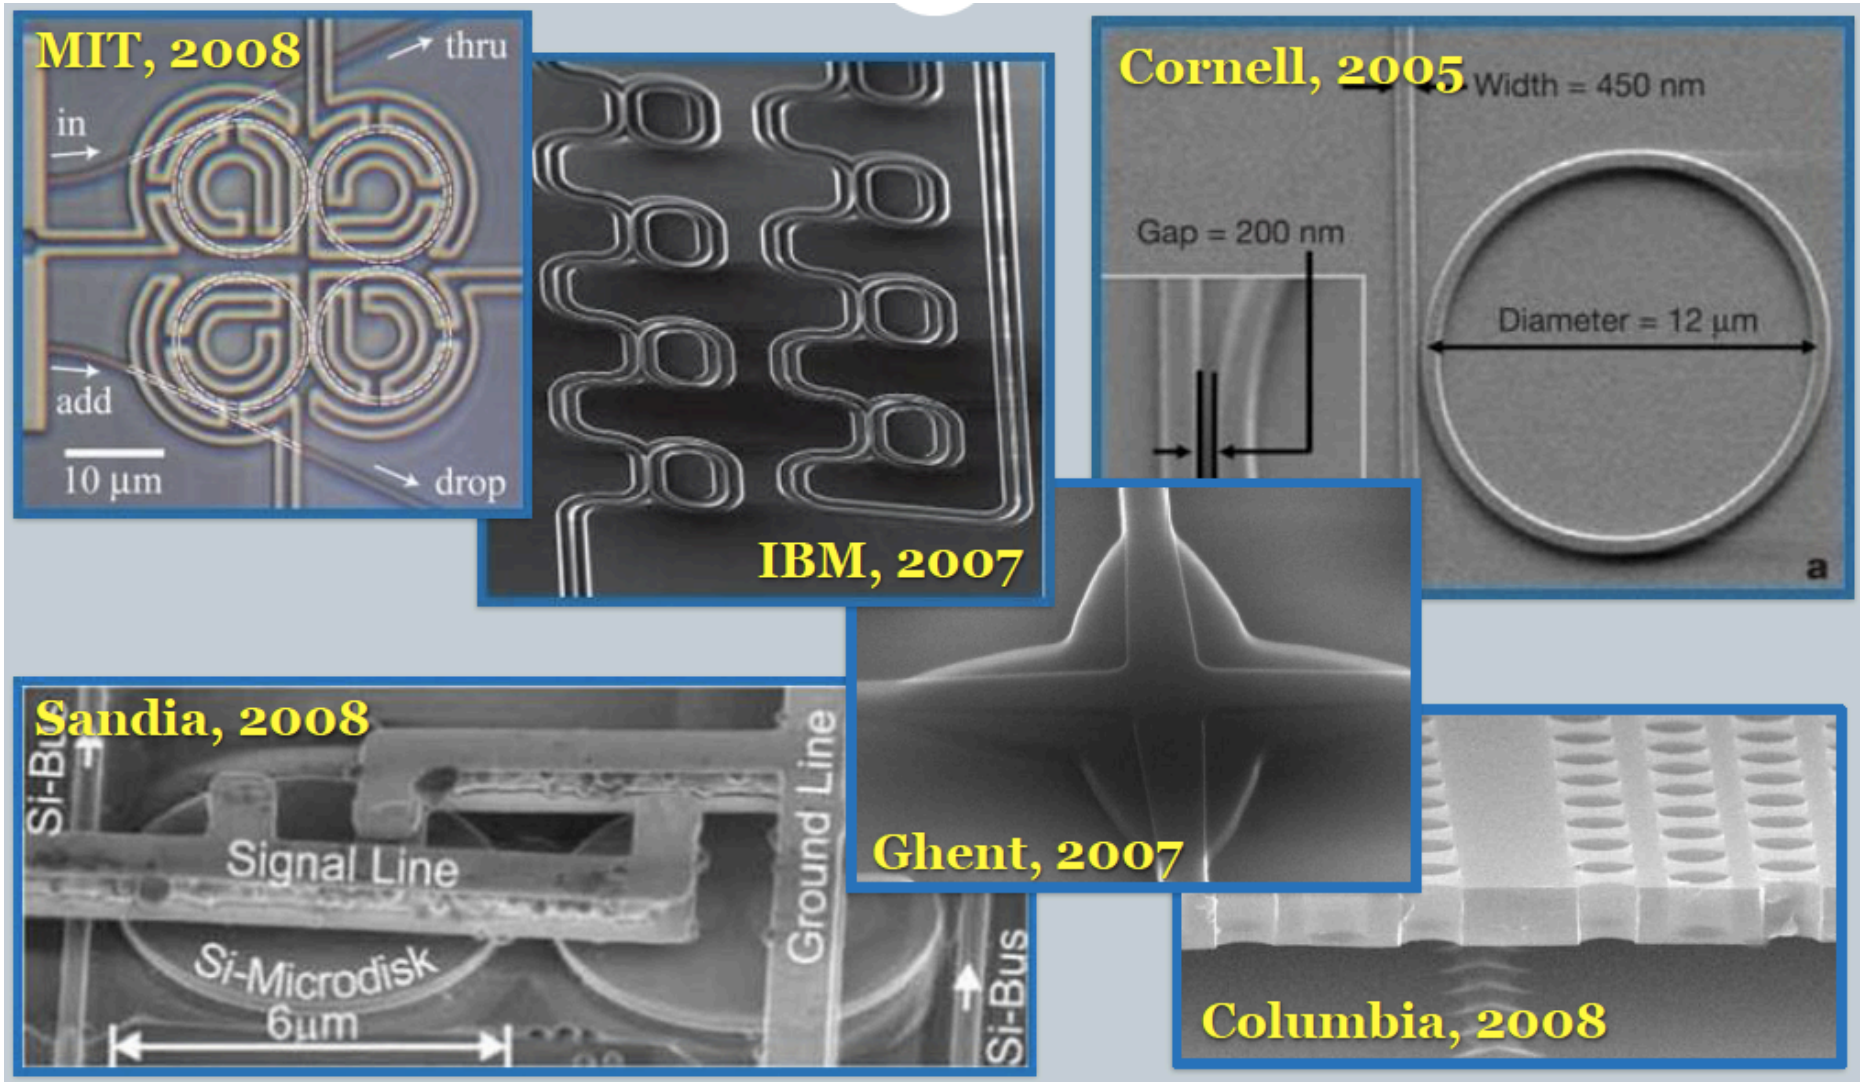

University of BRISTOL

First complete photonic link

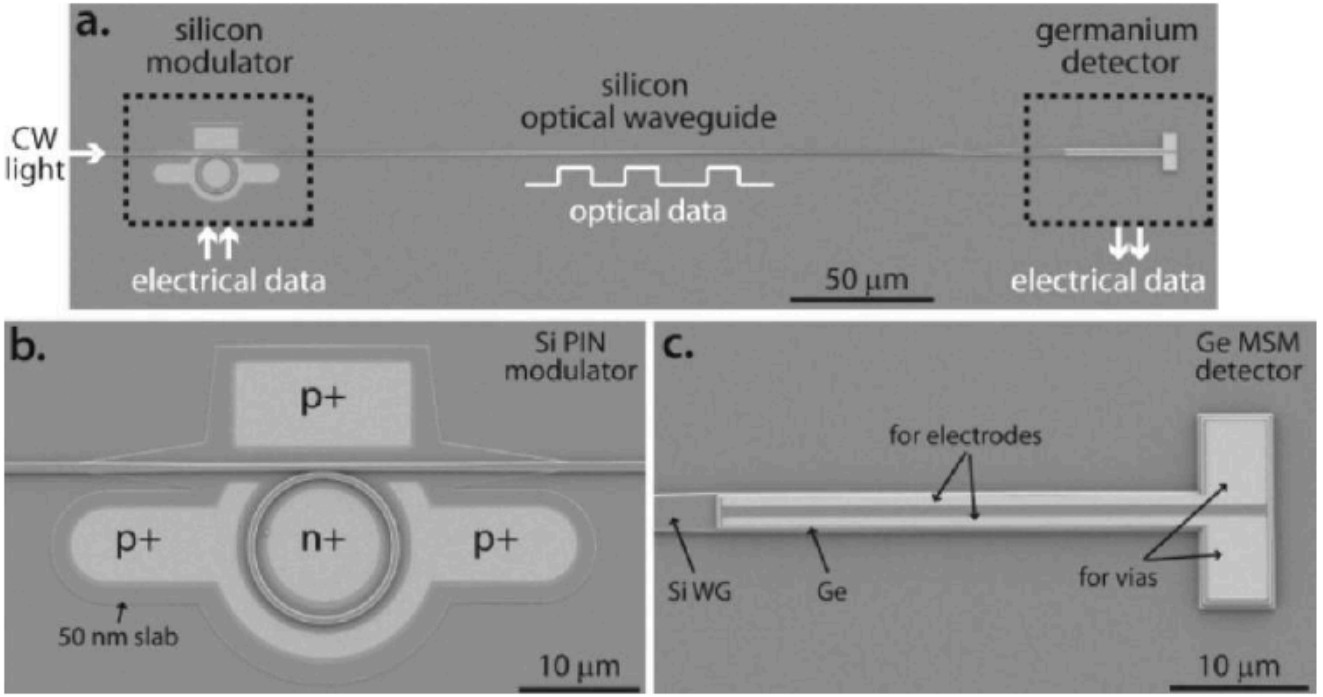

Integrated optical interconnect with silicon electro-optical modulator, silicon waveguide, and germanium-on-silicon photodetector

L. Chen, Optics Express, August 2009

University of BRISTOL

Does anyone else care about energy efficiency?

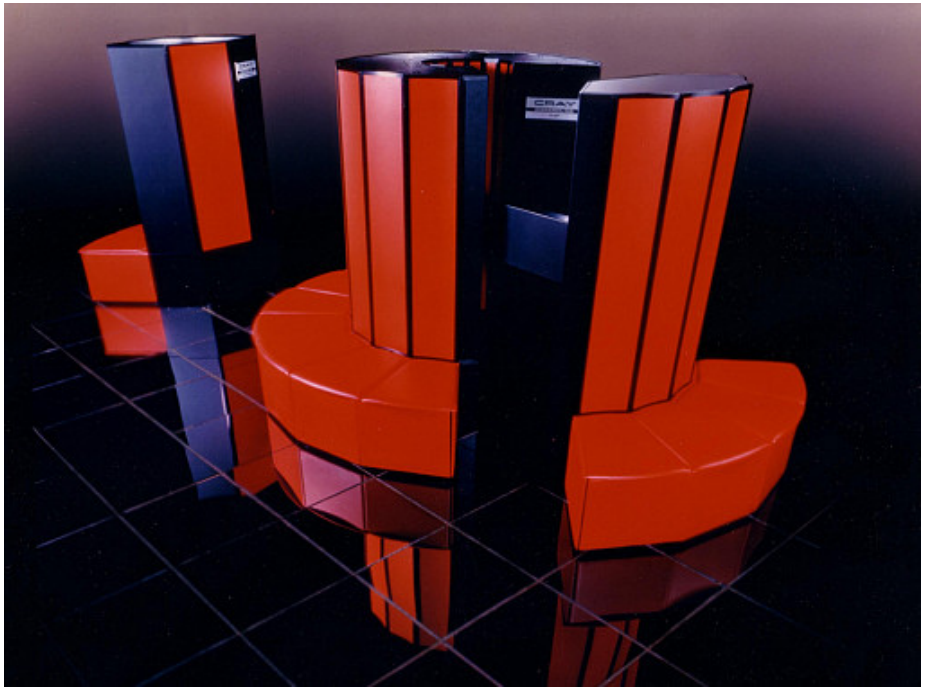

Cray X-MP, mid 1980's

Apple iPhone, 2007

University of BRISTOL

Mobile augmented reality

Layar.com

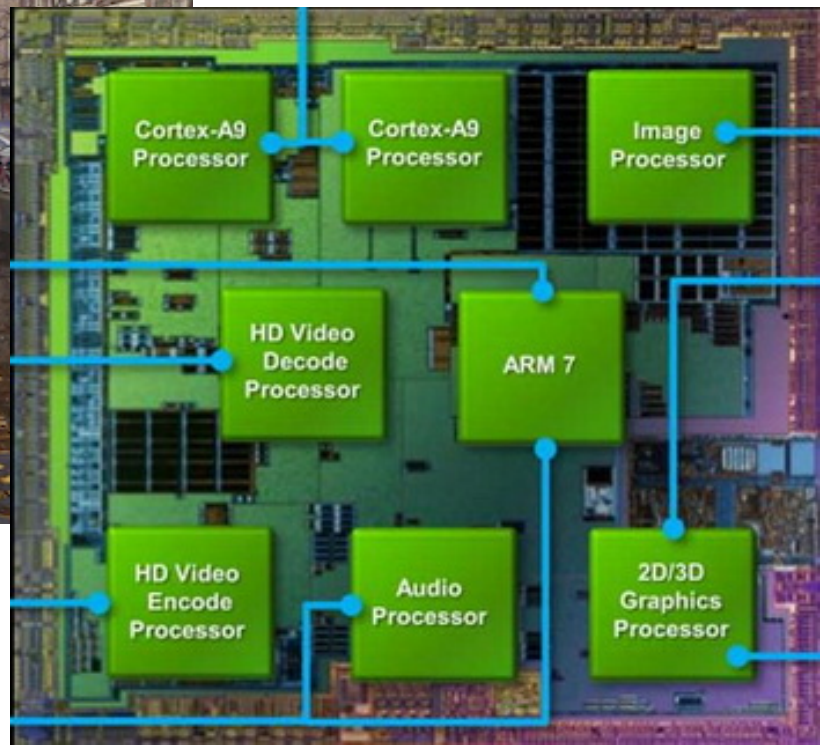

Nvidia Tegra 2

University of BRISTOL

Power efficiency in software

- It's not just hardware that can affect energy efficiency:
  - Moving data uses more energy than processing it
  - Bloated software uses more energy while sitting in dynamic memories (DRAMs)
  - It may be possible to improve performance per watt by 10-20% purely through changes to software optimisation

Bristol drug docking example

Therapy for Emphysema

Peptide libraries (based on a Trypsin inhibitor)

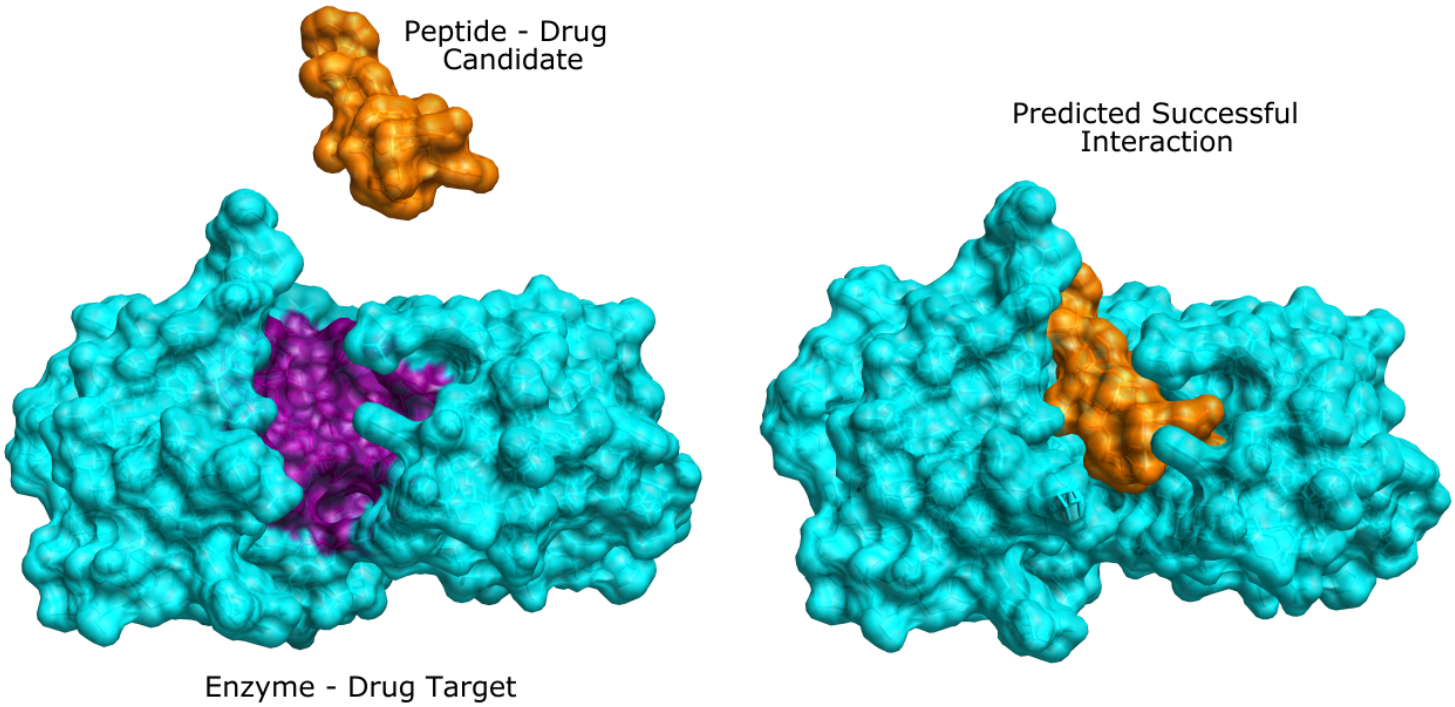

Flexible amino acid side-chains in both protein (receptor) and ligand (peptide)

University of BRISTOL

GPUs show energy efficiency

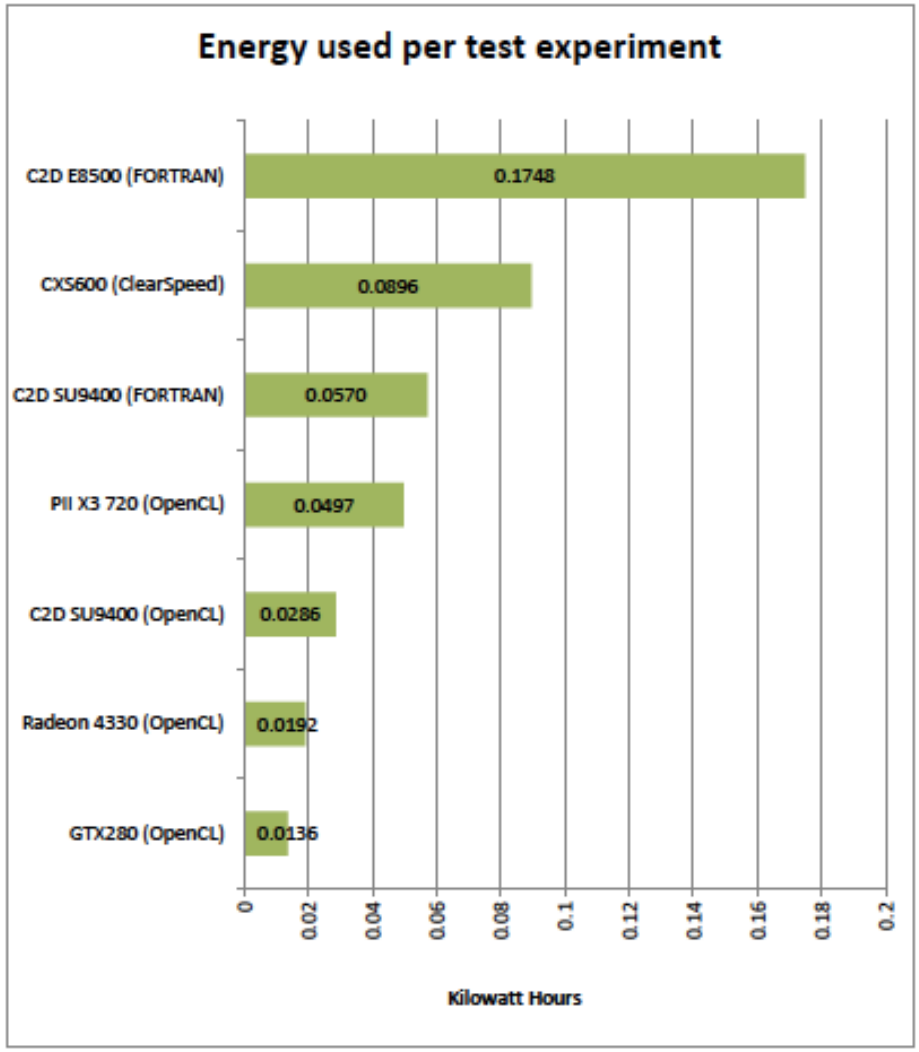

Dual core x86 CPU

AMD GPU

Nvidia C2050 GPU

12.8X difference between best and worst

University of BRISTOL

Optimising carbon emissions

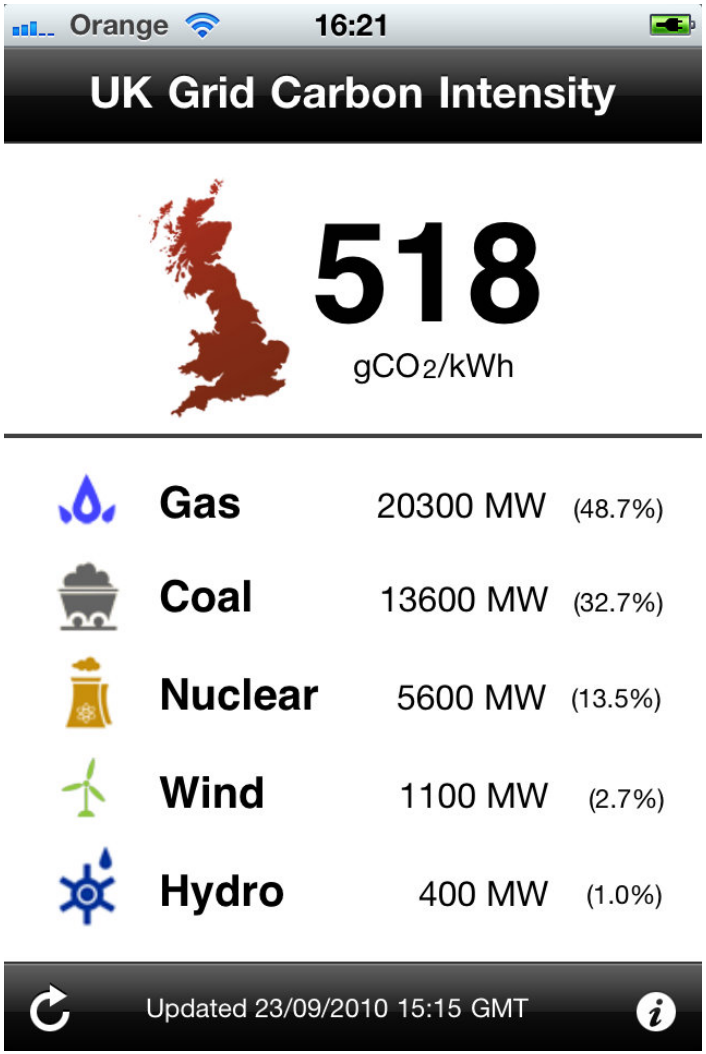

University of BRISTOL

Optimising carbon emissions

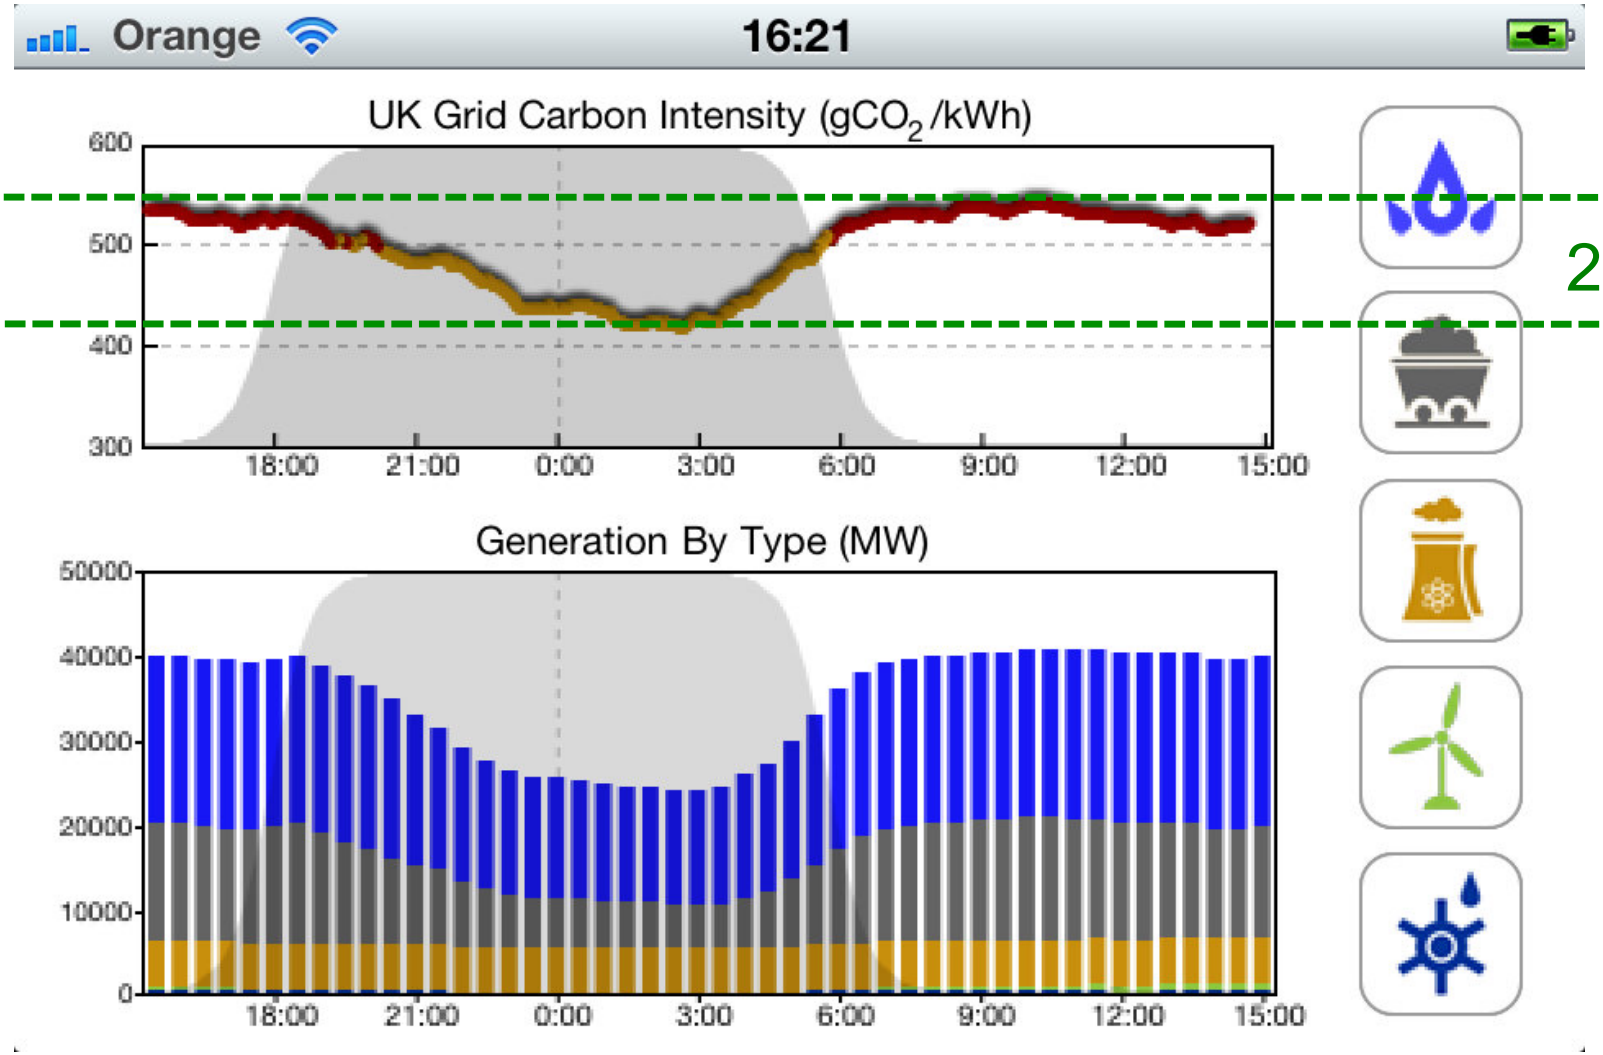

University of BRISTOL

Green 500

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	773.38	Forschungszentrum Juelich (FZJ)	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
1	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
1	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
4	492.64	National Supercomputing Centre in Shenzhen (NSCS)	Dawning Nebulae, TC3600 blade CB60-G2 cluster, Intel Xeon 5650/ nVidia C2050, Infiniband	2580
5	458.33	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband	276
5	458.33	IBM Poughkeepsie Benchmarking Center	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband	138
7	444.25	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband	2345.5
8	431.88	Institute of Process Engineering, Chinese Academy of Sciences	Mole-8.5 Cluster Xeon L5520 2.26 Ghz, nVidia Tesla, Infiniband	480
9	418.47	Mississippi State University	iDataPlex, Xeon X56xx 6C 2.8 GHz, Infiniband	72
10	397.56	Banking (M)	iDataPlex, Xeon X56xx 6C 2.66 GHz, Infiniband	72

* Performance data obtained from publicly available sources including TOP500

University of BRISTOL


ClearSpeed Visual Profiler
File
Window
Help
Buffer
598515025
Position
0
csvprof_trace.cst
Host Trace
Captured Metrics
MB/s
Degrees
(6261) FPGA BOARD TEMP [0]
(6261) FPGA CORE TEMP [0]
(6261) MTAP0 BOARD TEMP [0]
(6261) MTAP0 CORE TEMP [0]
(6261) MTAP1 BOARD TEMP [0]
(6261) MTAP1 CORE TEMP [0]
Degrees
80.0
77.5
75.0
72.5
70.0
67.5
65.0
62.5
60.0
57.5
55.0
52.5
50.0
47.5
45.0
0
50,000,000
100,000,000
150,000,000
200,000,000
250,000,000
300,000,000
350,000,000
400,000,000
450,000,000
500,000,000
550,000,000
usecs
Timeline
Buffer
usecs
0
61260500
122521000
183781500
245042000
306302500
367563000
428823500
490084000
551344500
141148851
(6261) MTAP1 CORE TEMP [0]
(6258) CSAPI_connect()
(6258) CSAPI_load()
(6258) CSAPI_allocate_static_shared_semaphore()
(6258) CSAPI_allocate_shared_memory()
(6258) CSAPI_run()
(6258) CSAPI_write_mono_memory()
(6264) CSAPI_write_mono_memory()
(6258) CSAPI_semaphore_signal()
(6258) CSAPI_semaphore_wait()
(6264) CSAPI_read_mono_memory()
(6258) CSAPI_wait_on_terminate()
(6258) CSAPI_free_memory()
(6258) Unknown
(6261) Unknown
(6264) Unknown
temperature = 58 c Degrees = 58
temperature = 75 c Degrees = 75
98391/98391 records loaded
200223282
200223283
1
200321835
341470686
141148851

Important takeaways

- Energy efficiency is becoming ***the*** first order consideration driving performance
- Heterogeneous computing is here to stay
- Even single chips will contain thousands of cores
- Hierarchies will become deeper
  - Processing, interconnect, memory, software
- Parallelism is now increasing exponentially and will continue to do so
- Software developers will have to start optimising their code for energy efficiency…