Major hardware trends affecting Exascale developments and their potential impact on software

Simon McIntosh-Smith simonm@cs.bris.ac.uk
Head of Microelectronics Research
University of Bristol, UK

University of BRISTOL

A brief biography

Graduated as Valedictorian in **Computer Science** from Cardiff University

1994 Joined **Inmos** to work for David May as a microprocessor architect

1999 Moved to **Pixelfusion** – a high-tech start-up designing the first many-core general purpose graphics processor (GPGPU)

2002 Co-founded **ClearSpeed** as Director of Architecture and Applications

2009 Head of Microelectronics Research at the **University of Bristol**, focusing on HPC and computer architecture. FP7 EESI member, PRACE prototype panelist, Archer UK national supercomputer project group

Microelectronics Research in Bristol

Simon McIntosh-Smith
Head of Group

Prof David May

Prof Dhiraj Pradhan

Dr Jose Nunez-Yanez

Dr Kerstin Eder

Dr Simon Hollis

Dr Dinesh Pamunuwa

7 tenured staff, 6 research assistants, 16 PhD students

University of BRISTOL

Bristol Research

Energy Aware COmputing (EACO):

- **Multi-core and many-core computer architectures** (FP7 EESI)
  - ClearSpeed, XMOS, Inmos, Pixelfusion, …
- **Algorithms for *heterogeneous architectures***
  - CPUs+GPUs, OpenCL
- **Electronic and Optical Network on Chip** (NoC)
- **Fault tolerant design** (hardware and software)
  - Near threshold computing for embedded medical devices (FP7 DeSyRe)
- Reconfigurable architectures (FPGA)
- Design verification (formal and simulation-based), formal specification and analysis
- Silicon process variation
- Design methodologies, modelling & simulation of MNT based structures and systems

Context

University of BRISTOL

Moore's Law today

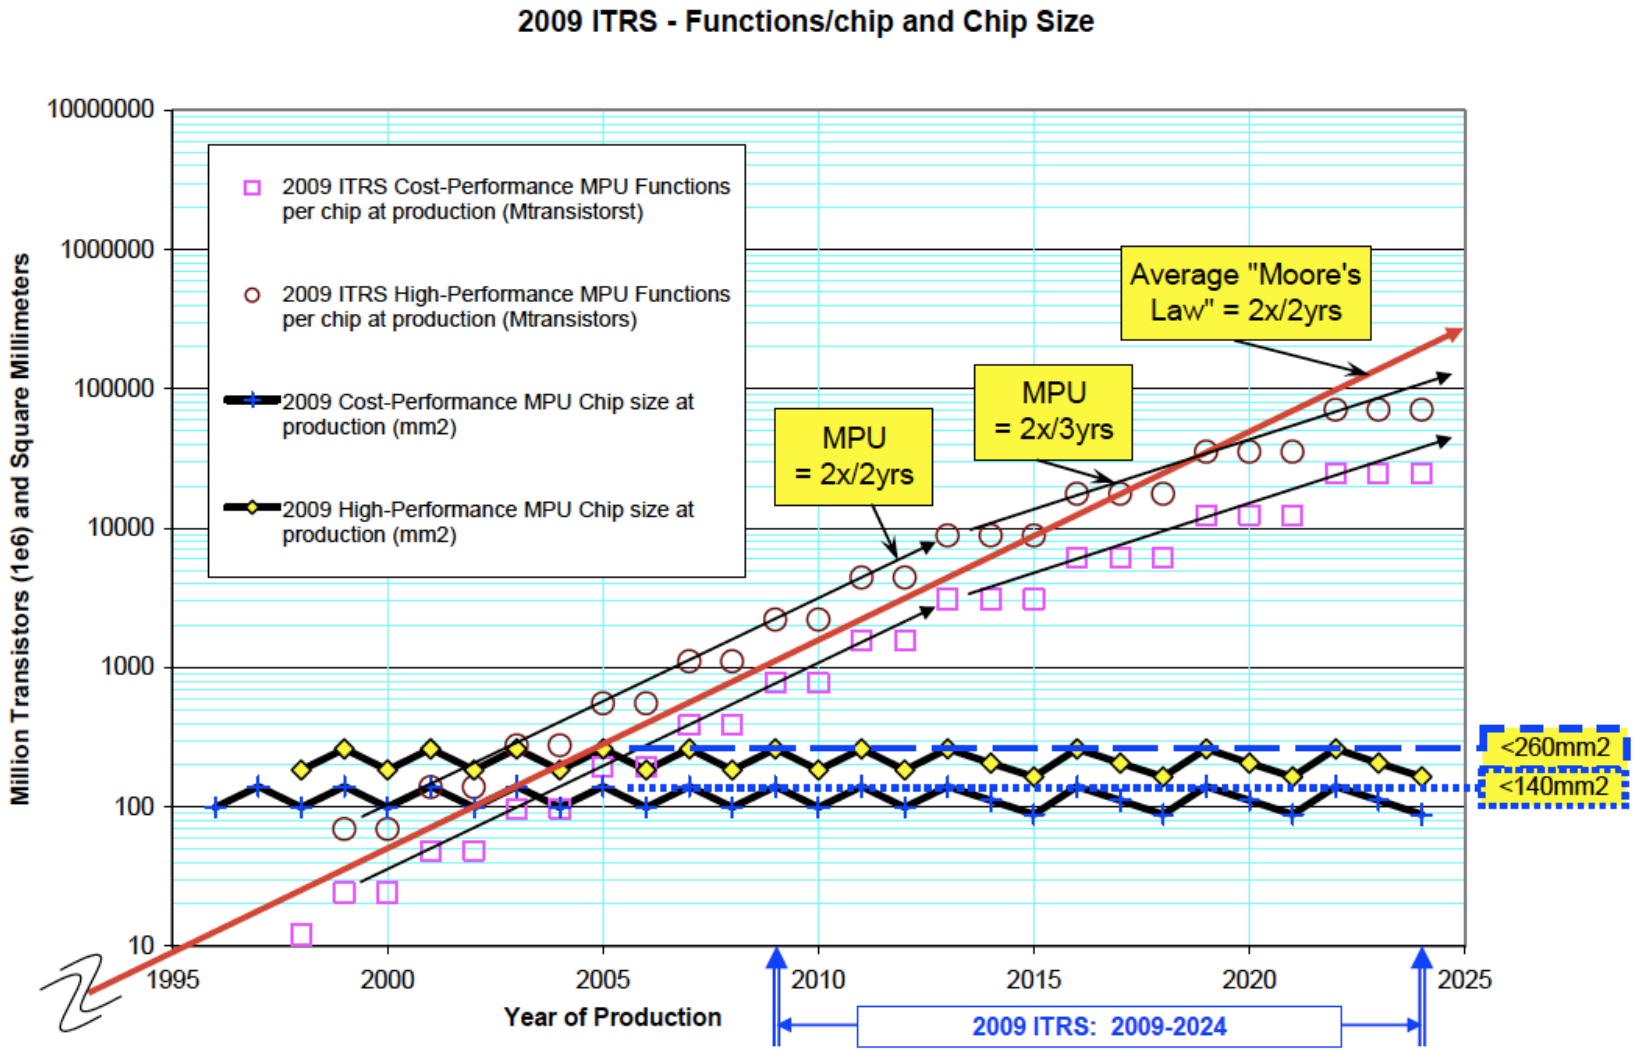

http://www.itrs.net/Links/2009ITRS/2009Chapters_2009Tables/2009_ExecSum.pdf

Moore's Law today

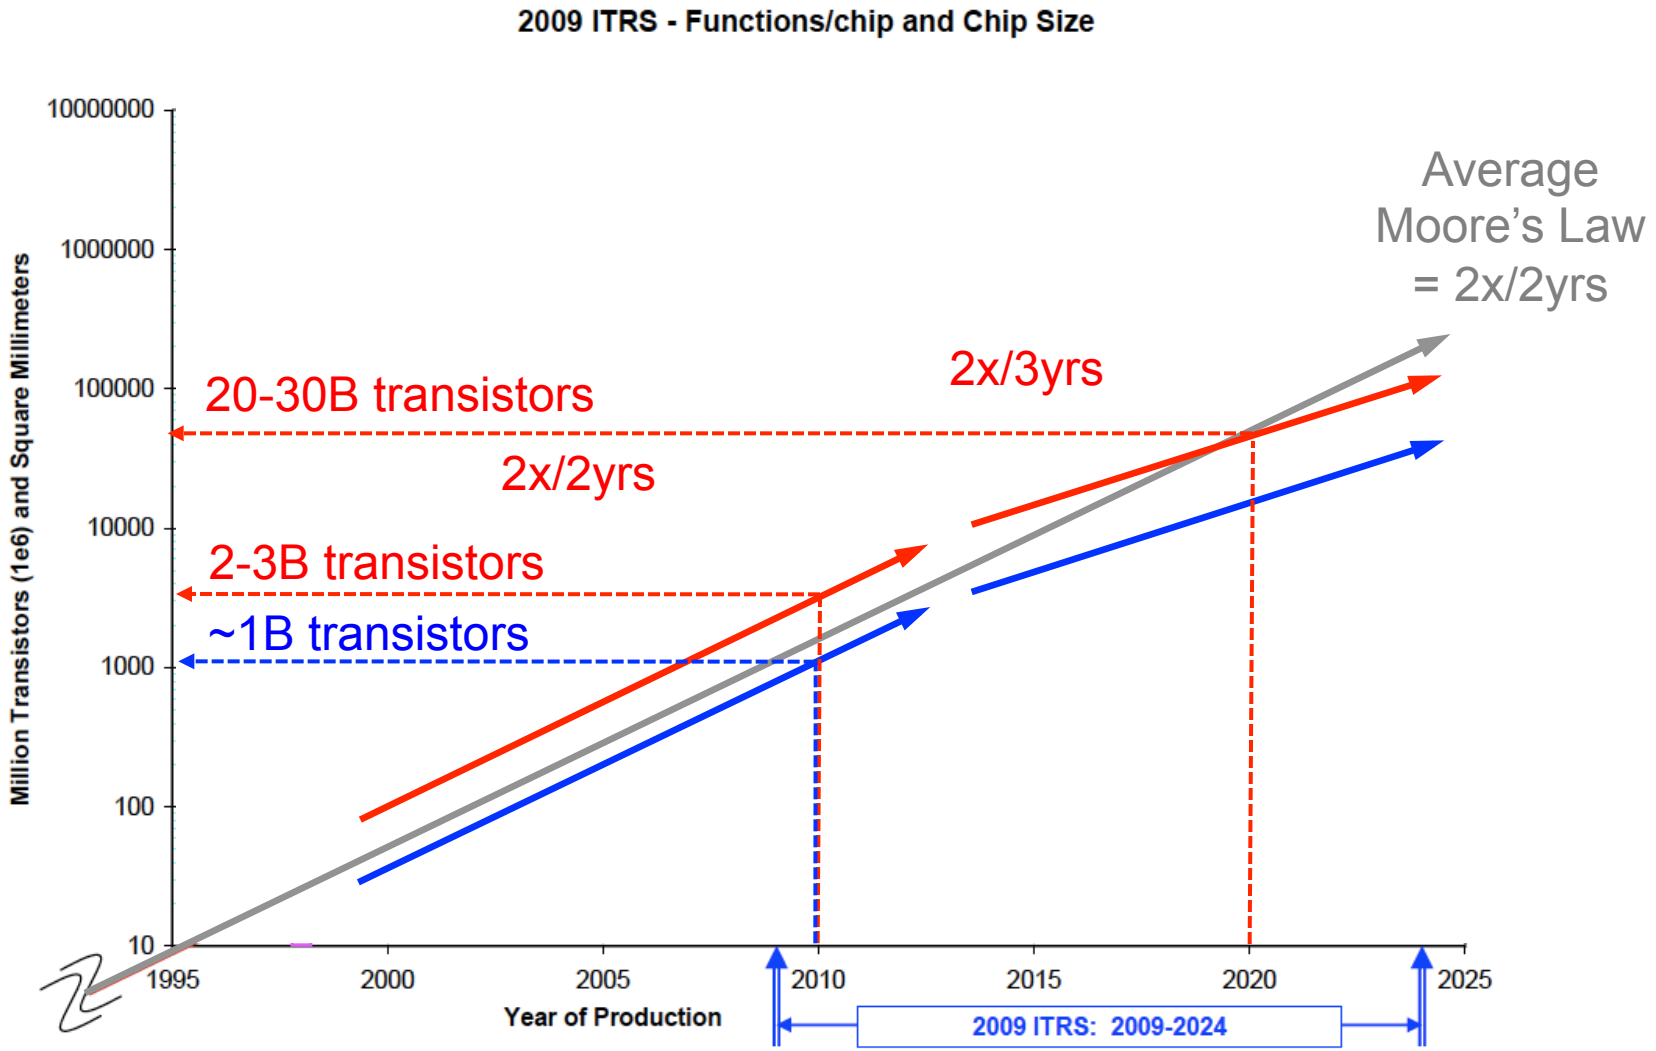

University of BRISTOL

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

Important technology trends

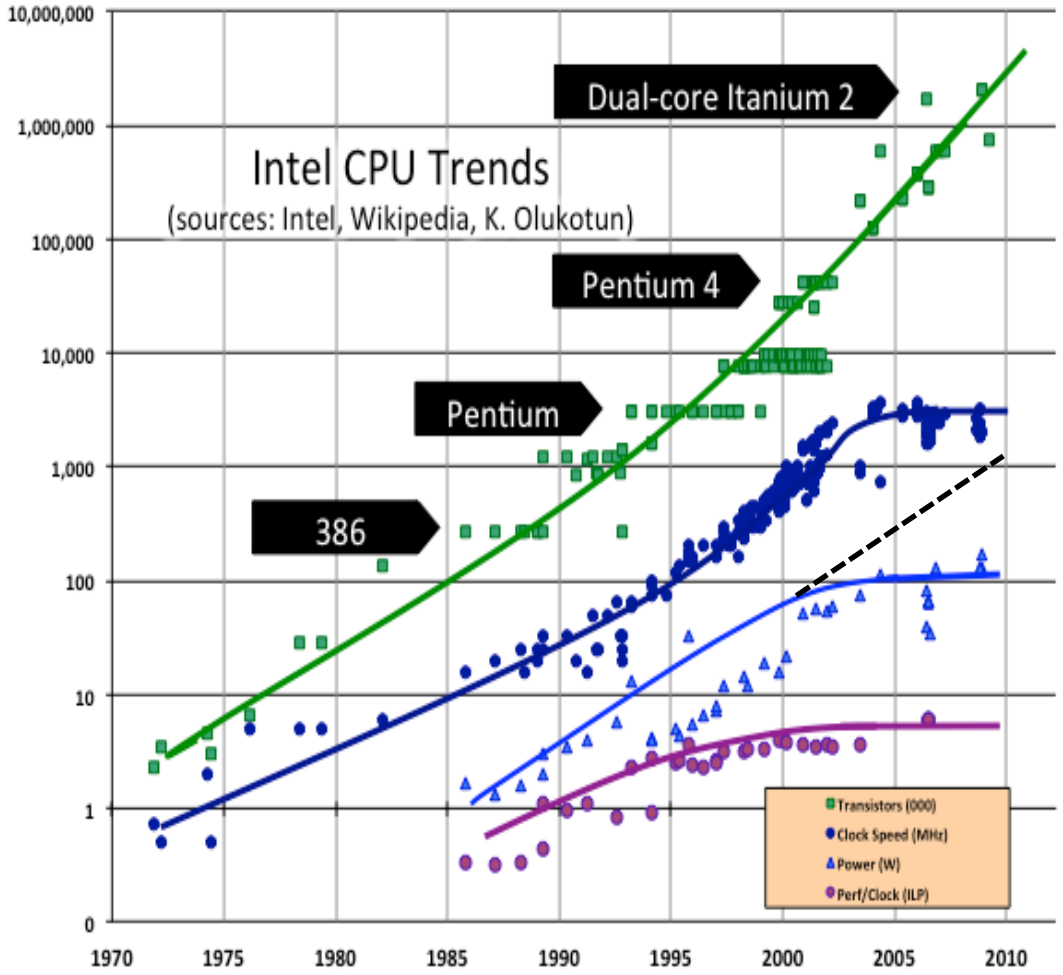

The real Moore's Law

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

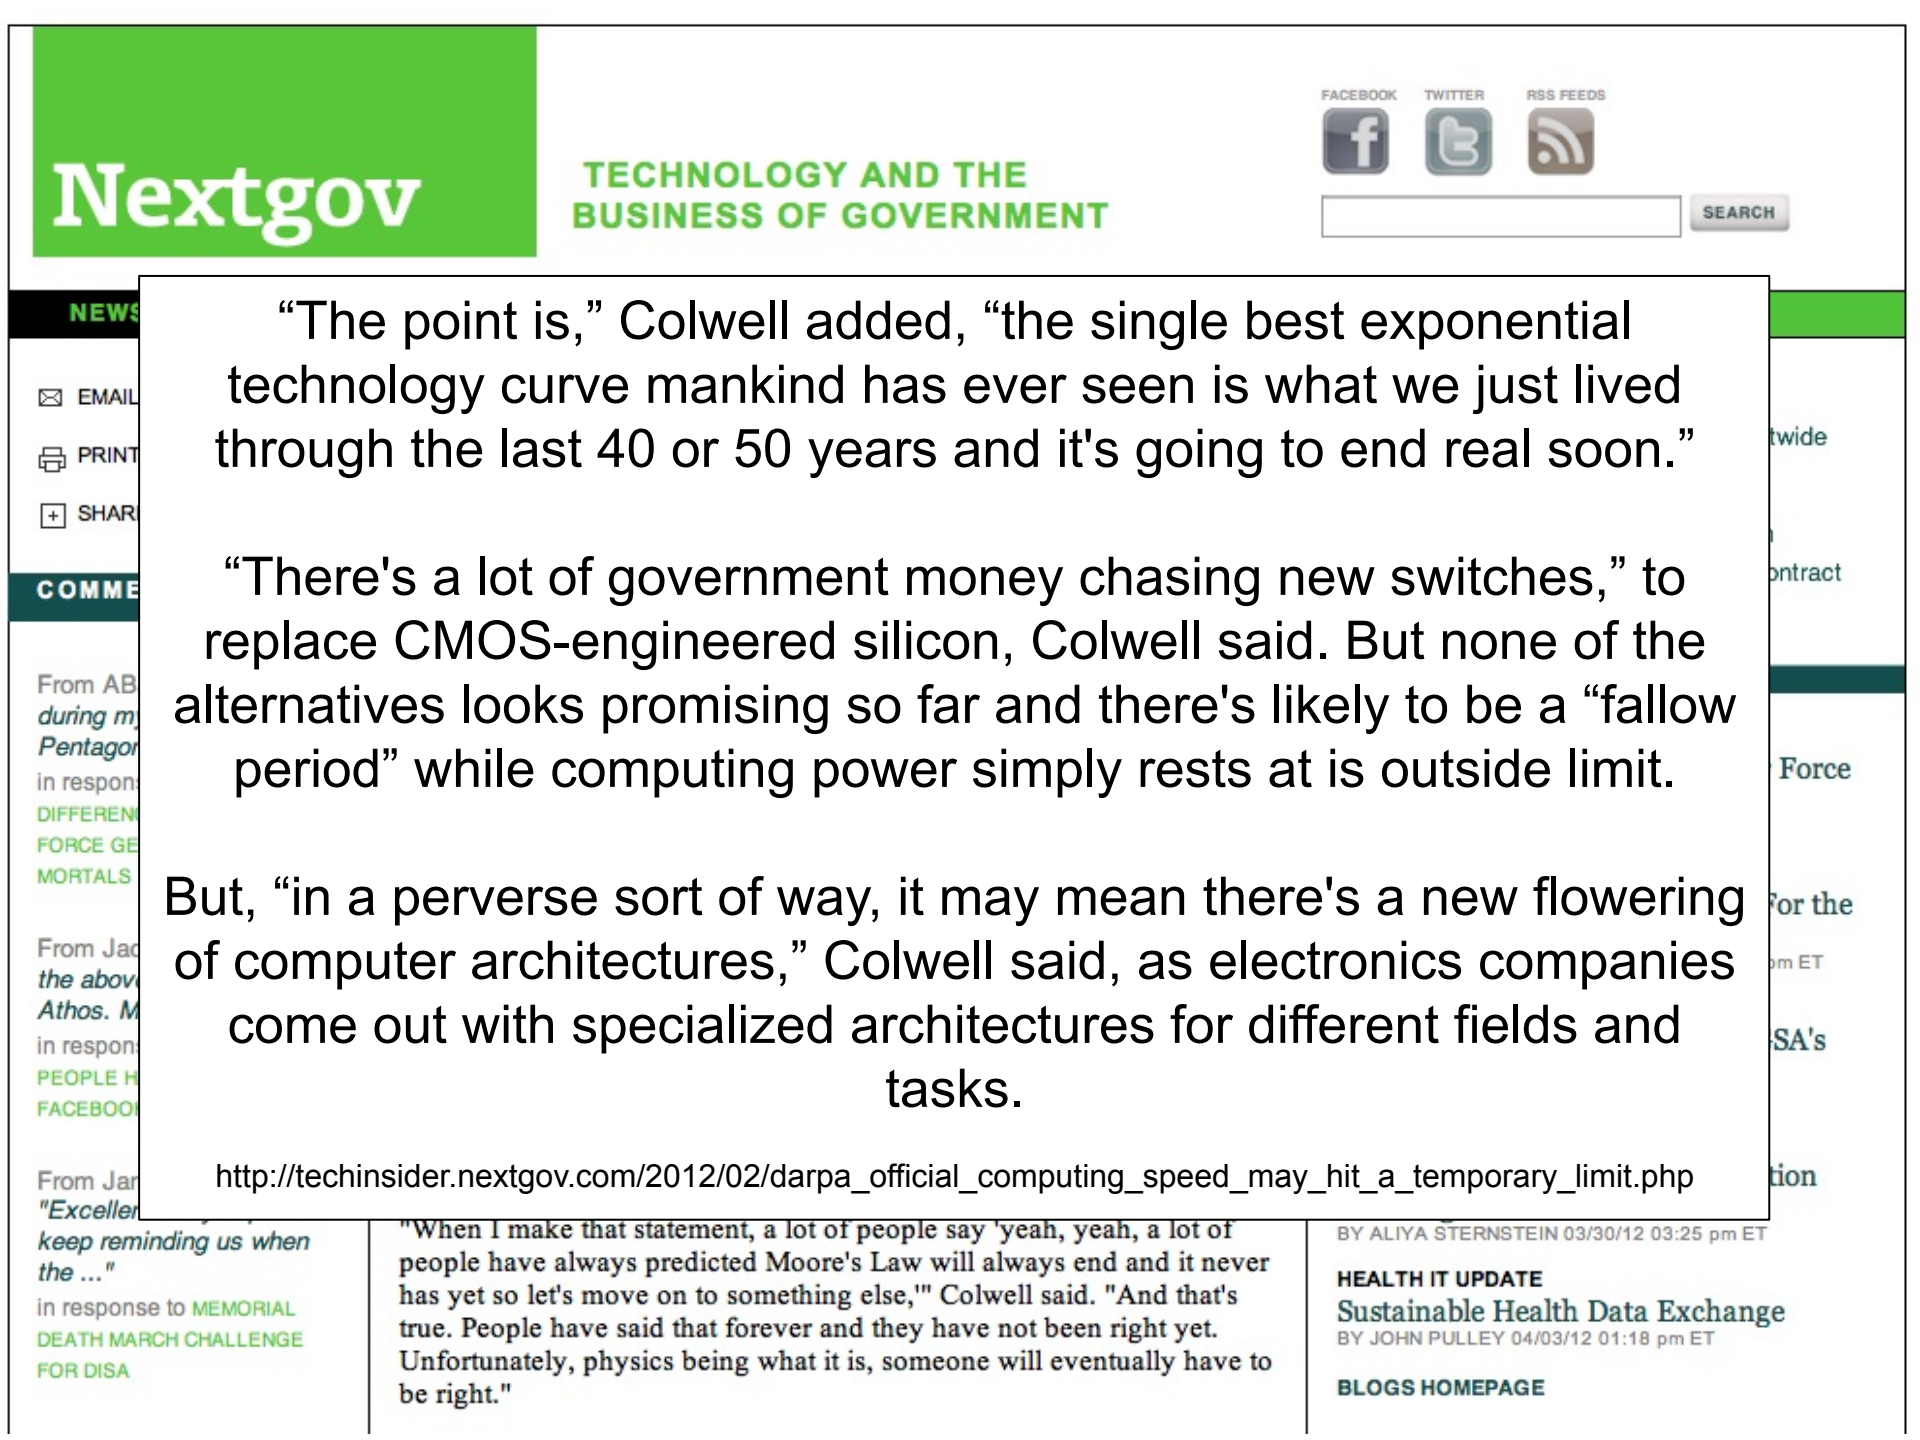

"The point is," Colwell added, "the single best exponential technology curve mankind has ever seen is what we just lived through the last 40 or 50 years and it's going to end real soon."

"There's a lot of government money chasing new switches," to replace CMOS-engineered silicon, Colwell said. But none of the alternatives looks promising so far and there's likely to be a "fallow period" while computing power simply rests at is outside limit.

But, "in a perverse sort of way, it may mean there's a new flowering of computer architectures," Colwell said, as electronics companies come out with specialized architectures for different fields and tasks.

http://techinsider.nextgov.com/2012/02/darpa_official_computing_speed_may_hit_a_temporary_limit.php

Major hardware trends

University of BRISTOL

The five major hardware trends that will affect exascale software

1. Changes to memory hierarchies
2. The impact of fault tolerance
3. Focus on energy efficiency
4. Heterogeneity
5. Scale

Changes to memory hierarchies

University of BRISTOL

Changes to memory hierarchies

Causes:

- 3D chip stacking
- Photonic interconnects
- Continued deepening of memory hierarchies

3D stacked memories

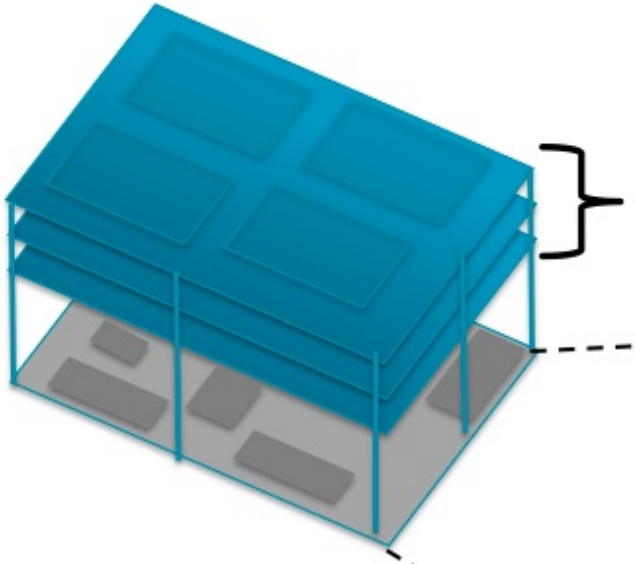

3D-stacked
DRAM cache

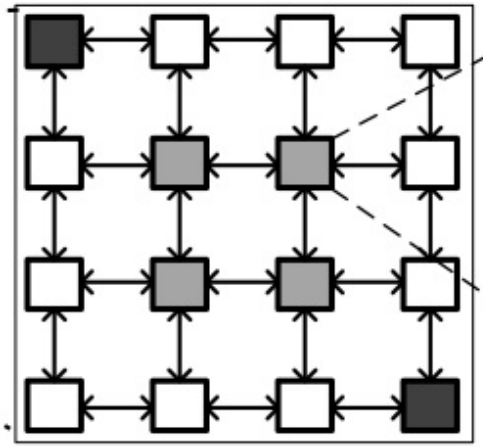

Logic die

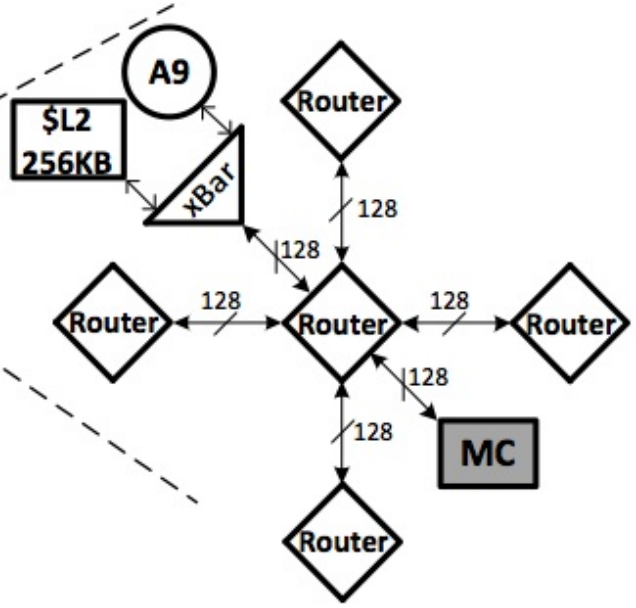

Node

- 16 A9 ARM cores
- 4MB shared L2 cache
- 4x4 mesh interconnect
- 4 memory controllers with Wide I/O interface
- 4GB on-chip DRAM

3D stacked memories

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

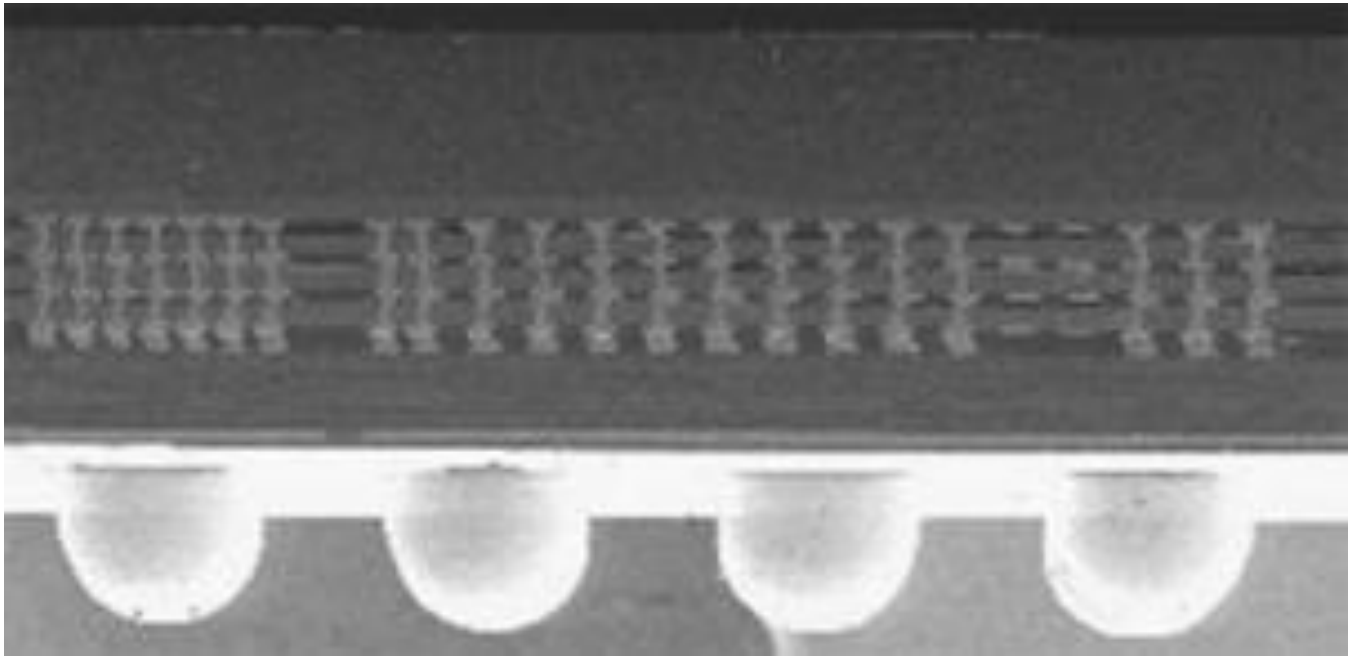

Samsung 3D DRAM

Photonic networks

- Roadmaps to achieve ExaFLOPs (10^{18}) by 2018 are relying on some major hardware breakthroughs to improve energy efficiency
- Prof Jeremy O'Brien, Centre for Quantum Photonics, University of Bristol, also Prof Keren Bergmen's work at Columbia
- ***Moving data*** becoming an increasingly dominant fraction of energy dissipation in microelectronics
  - "Compute free, bandwidth expensive"

Photonic NoC integration

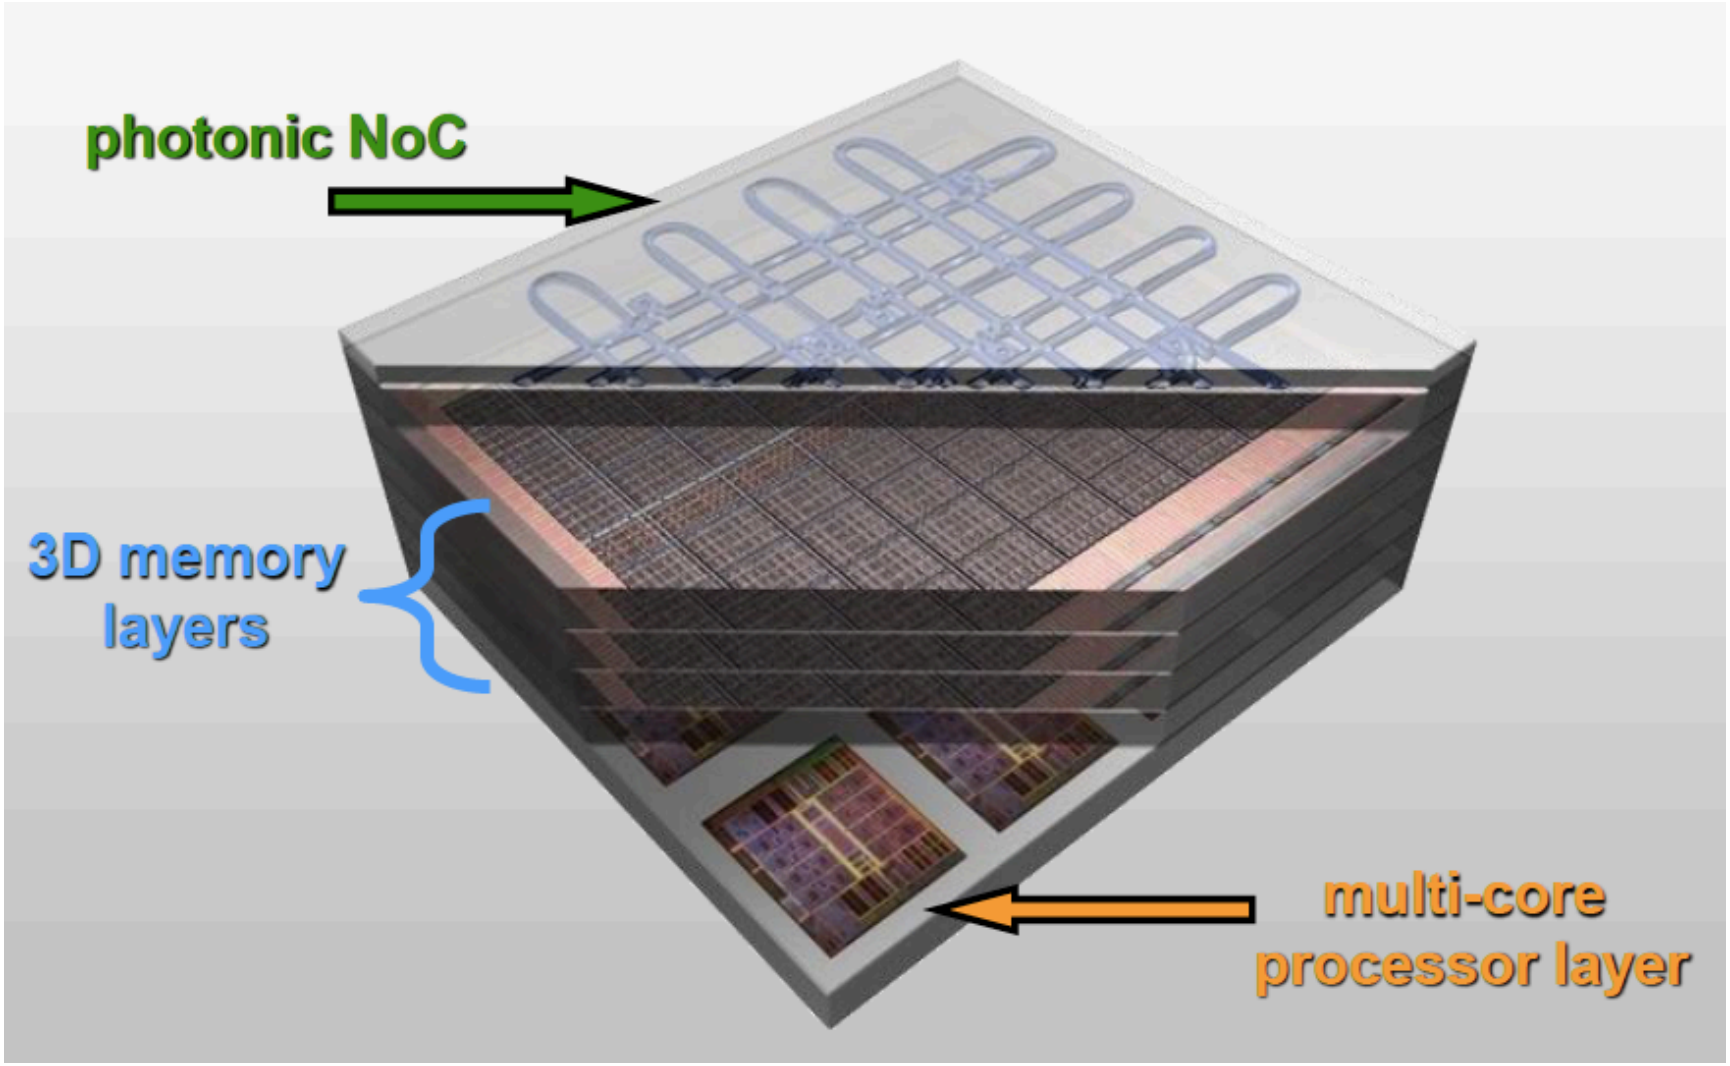

Optically interconnected supercomputing board

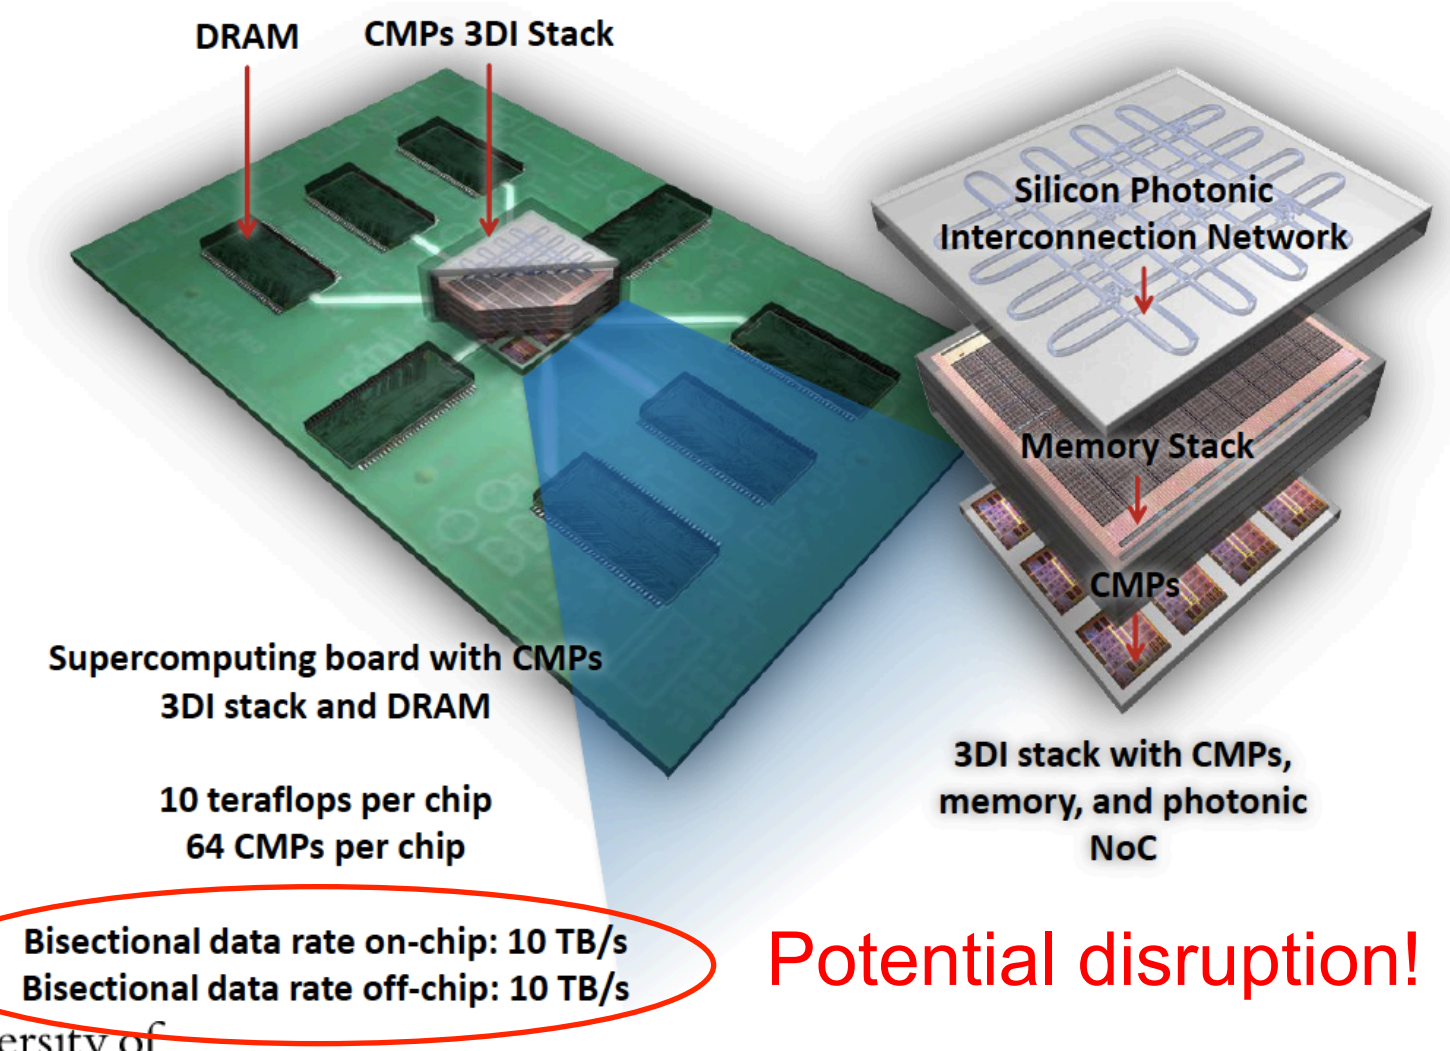

Potential disruption!

Optical computing in Bristol

Coupled waveguides

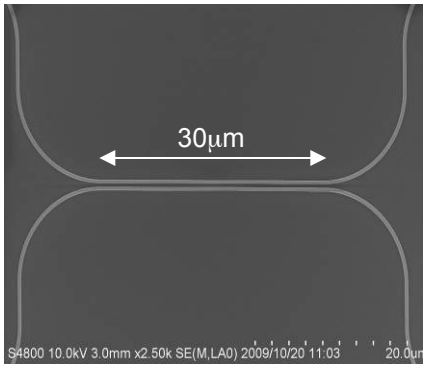

Splitter

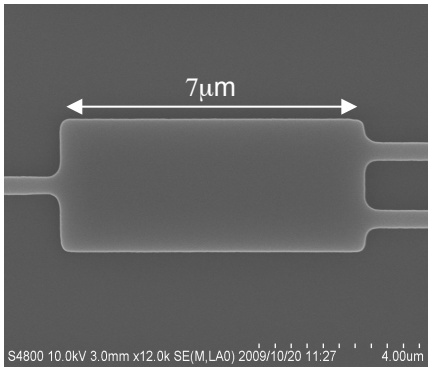

Ring resonator

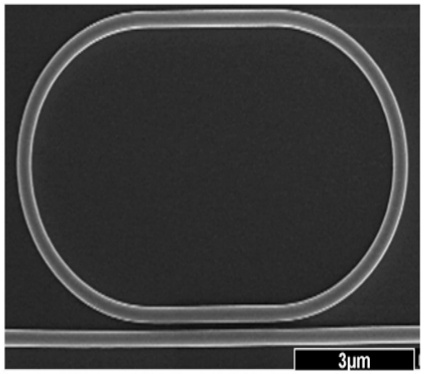

Spiral delay line

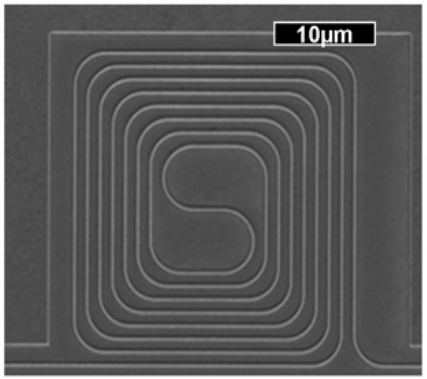

http://www.phy.bris.ac.uk/groups/cqp/index.html

University of BRISTOL

Centre for Quantum Photonics

Silicon photonic components

2x2 splitter

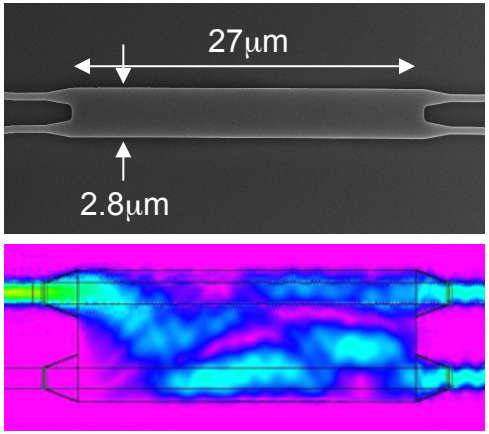

Wavelength filtering

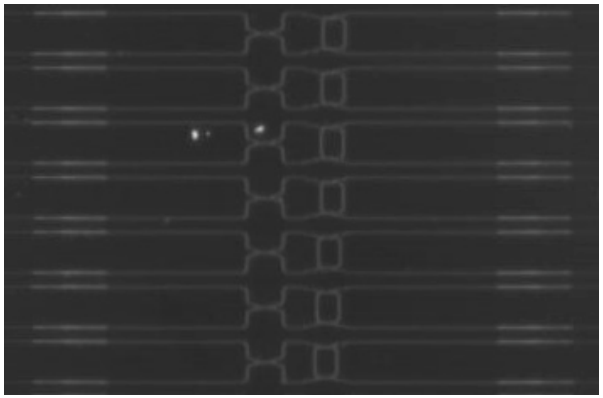

16 mode coupler

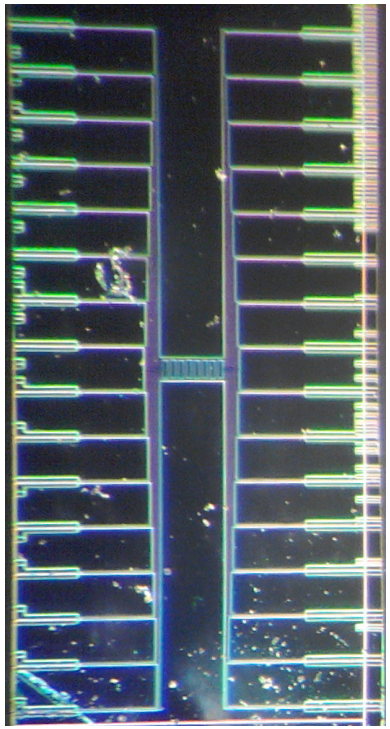

1x4 dynamic router

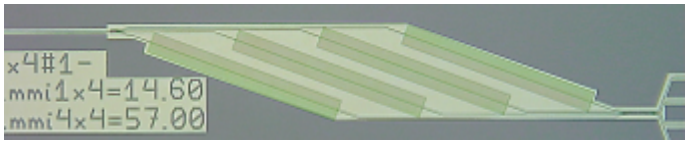

Si waveguide

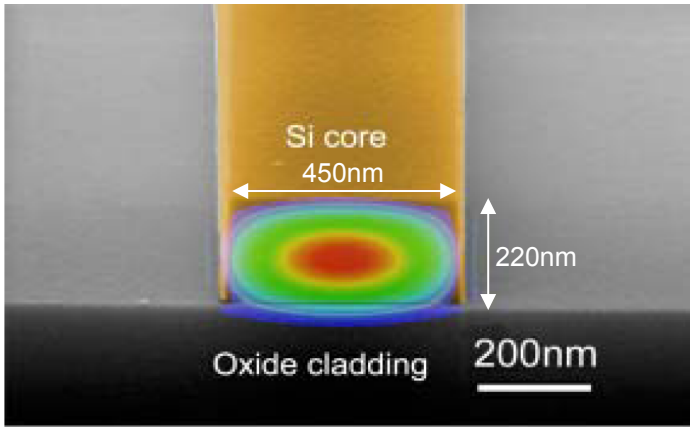

1x4 splitter

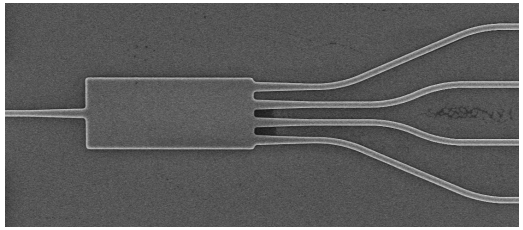

University of BRISTOL

4 Port Optical Router

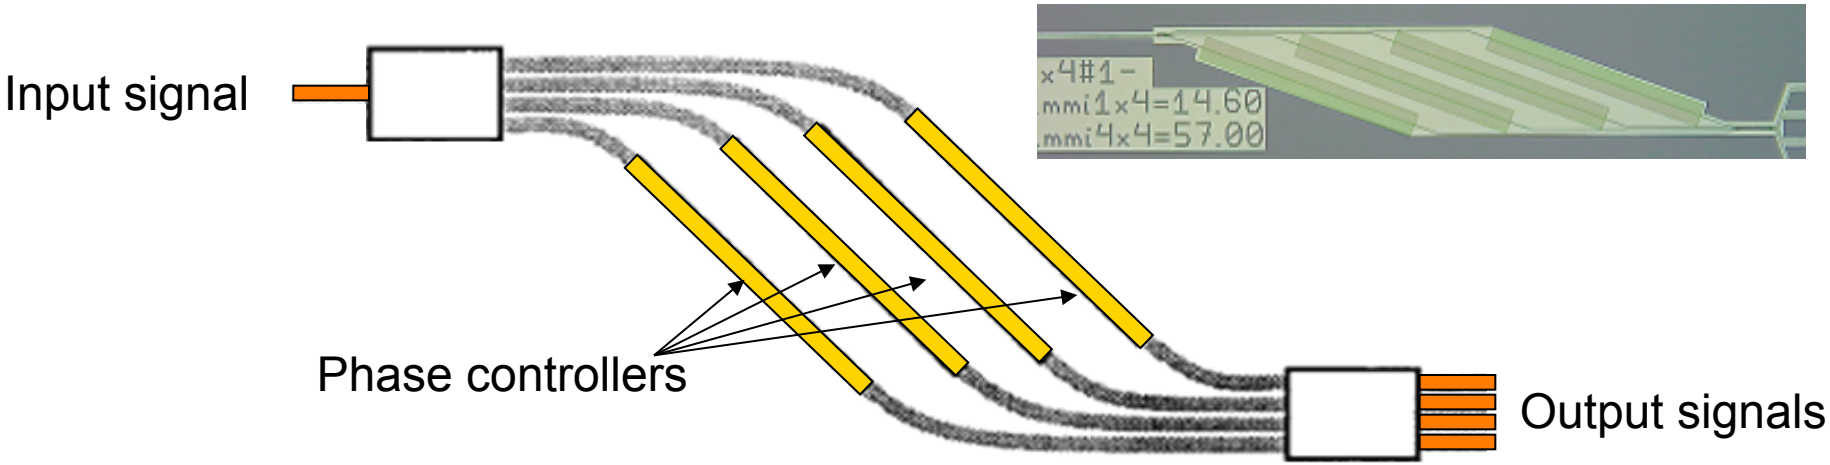

- Switch and route light to different parts of the optical network
- Phase controllers set where the light goes

http://www.phy.bris.ac.uk/groups/cqp/index.html

University of BRISTOL

Centre for Quantum Photonics

Implications

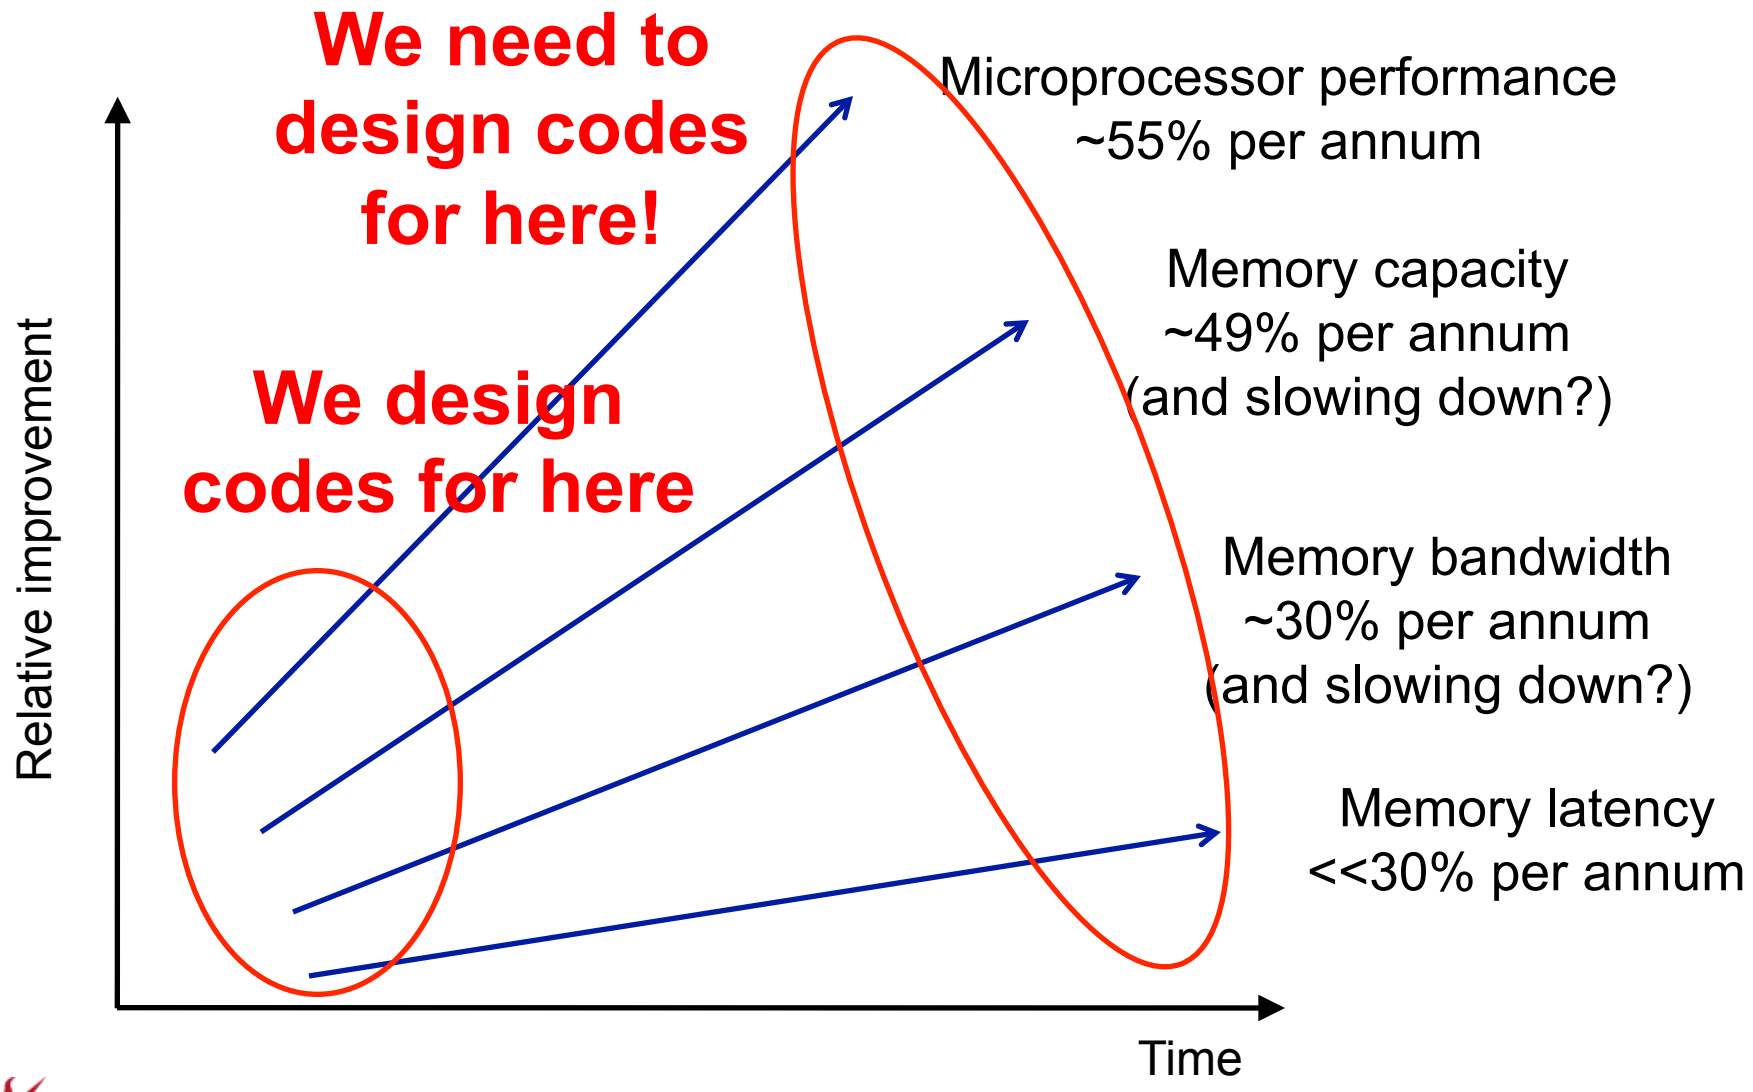

The impact of fault tolerance

University of BRISTOL

The impact of fault tolerance

Causes:

- Shrinking feature sizes
- Near-threshold operating voltages

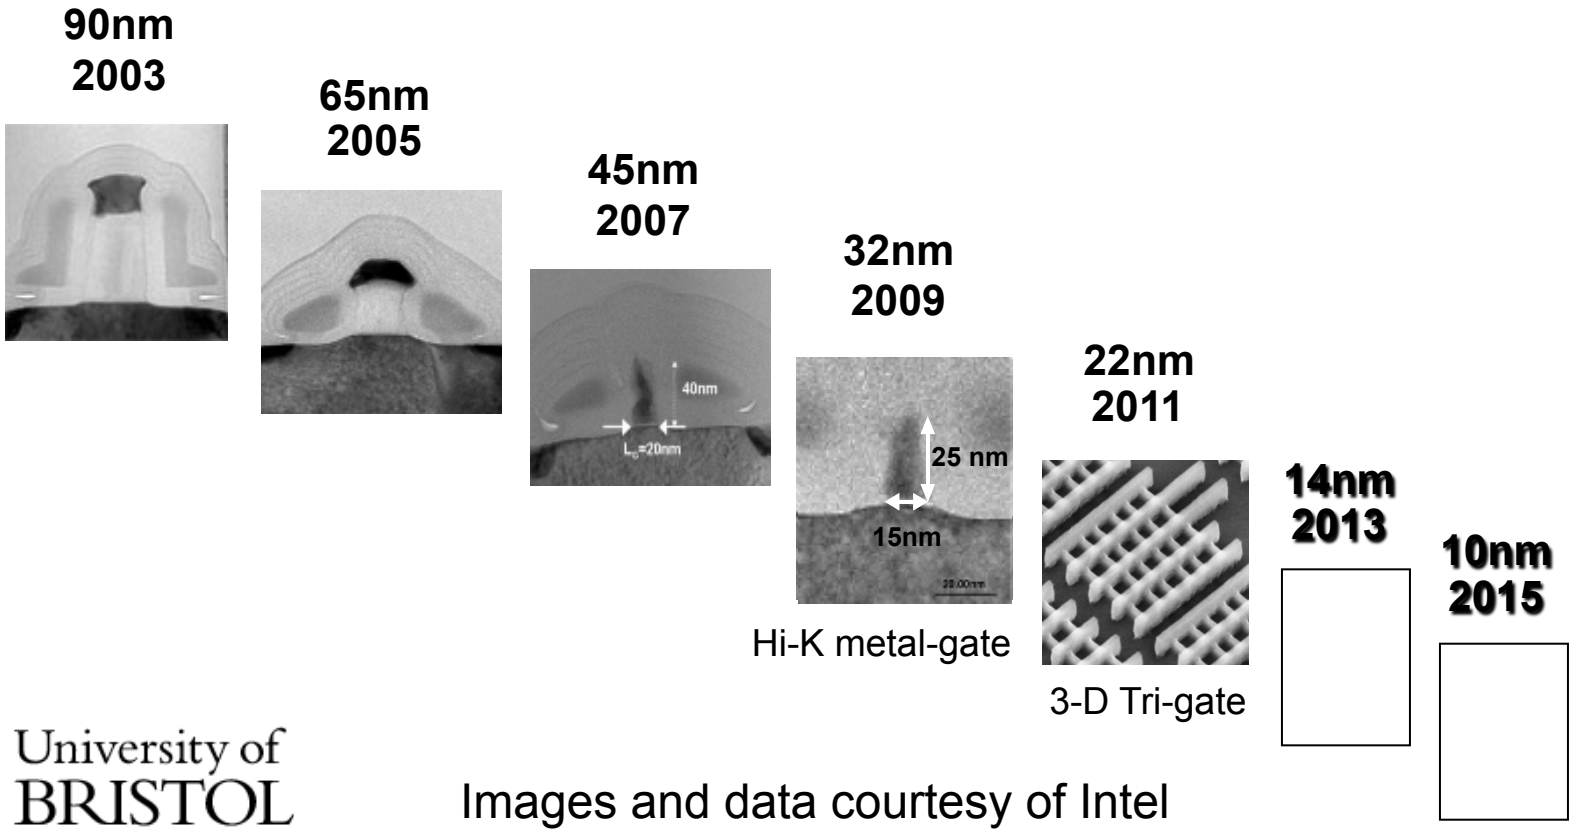

Images and data courtesy of Intel

Implications

- I believe high-end systems vendors ***will*** be able to provide reliable Exascale hardware, ***<u>at a price</u>***
  - Significant energy and silicon cost
- Fault tolerant software (applications, middleware, OS, …) will open up other possibilities
- Great need for tools and techniques to make software developers “fault aware”

Energy Efficiency Changes

University of BRISTOL

Focus on energy efficiency

Hardware energy efficiency improvements alone won't be enough.

Significant energy efficiency improvements could come from software optimisations:

- Tradeoff data movement vs. computation
- Mixed precision
- Asynchronous/event driven vs. polling
- And many, many more…

EESI-2 FP7 project to include EEHPC focus

Electronics Weekly 15-21 June 2011

news

LOW POWER

Lack of software support marks the low power scorecard at DAC

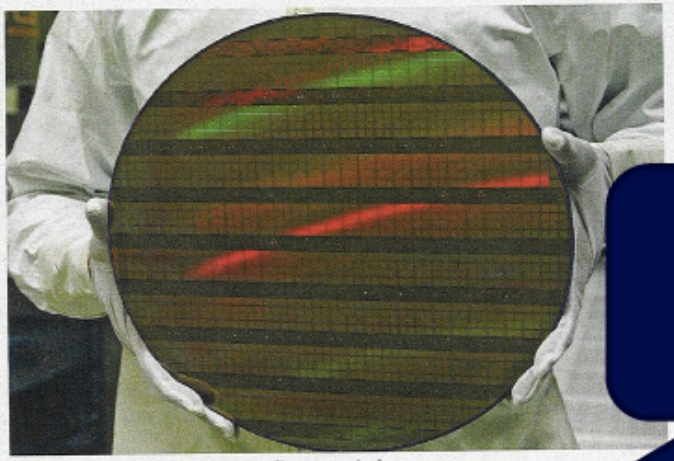

Intel waits for better low-power software control

power circuitry could save maybe 20% in a typical multimedia-oriented core.

Make the software controlling it better at controlling the power states and that difference could be three to five times.

During an afternoon panel discussion Ambrose Low, director of design engineering at Broadcom said: "We have hundreds of knobs in the hardware to turn power down.

"The question is whether we can take the actual use-cases into consideration and optimise the software to

...they can save," he added.

Chris Edwards writes the Low-Power Design Blog (enabled by Mentor Graphics) on ElectronicsWeekly.com

www.electronicsweekly.com/ew-blogs/

[...] if the software keeps cores active for no good reason [the hardware] won't deliver a realised saving.

With limited software support, dedicated low-power circuitry could save maybe 20%

Make the software better at controlling the power states and that difference could be **three to five times.**

University of BRISTOL

Optimising use of multiple cores

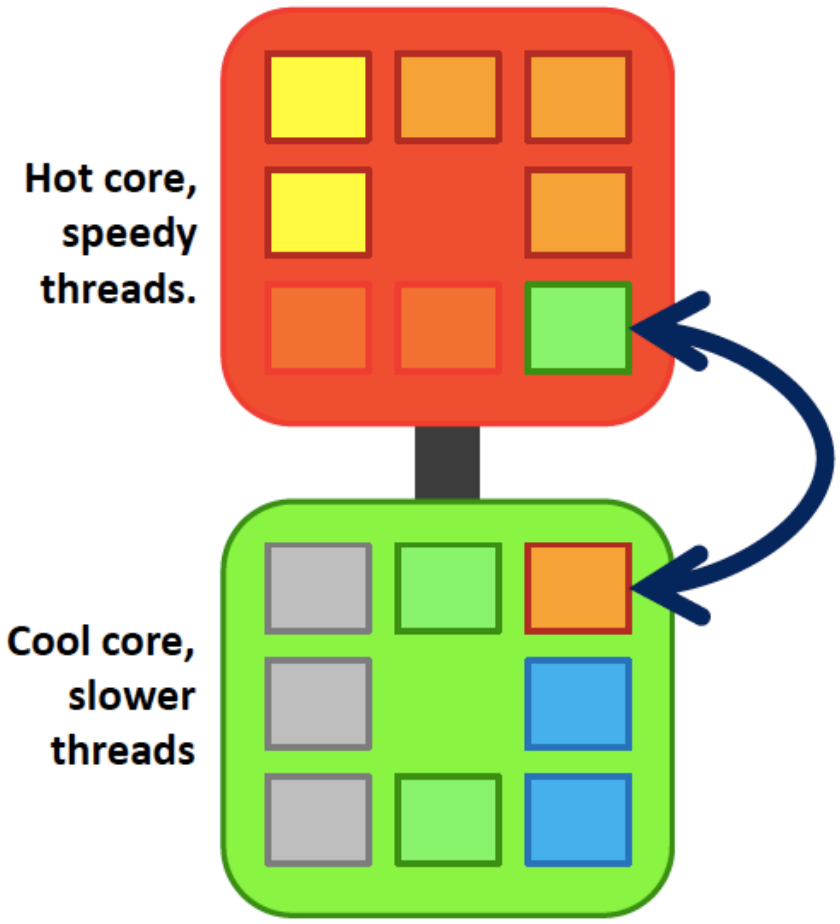

- Assign less demanding tasks to low frequency cores
- Move tasks that interfere with other optimisation efforts onto a different core
- Example: ARM's big.LITTLE SoC design strategy
- We have to deal with the timing & communication implications of doing this.

"Energy Efficient HPC Software: Learning from Embedded Systems Development", K. Eder, S. Kerrison and S. McIntosh-Smith, SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP12), Savanah, Georgia, February 2012 http://www.cs.bris.ac.uk/~simonm/publications/sk_siampp12_eehpcs.pdf

Embedded processors: the next disruption?

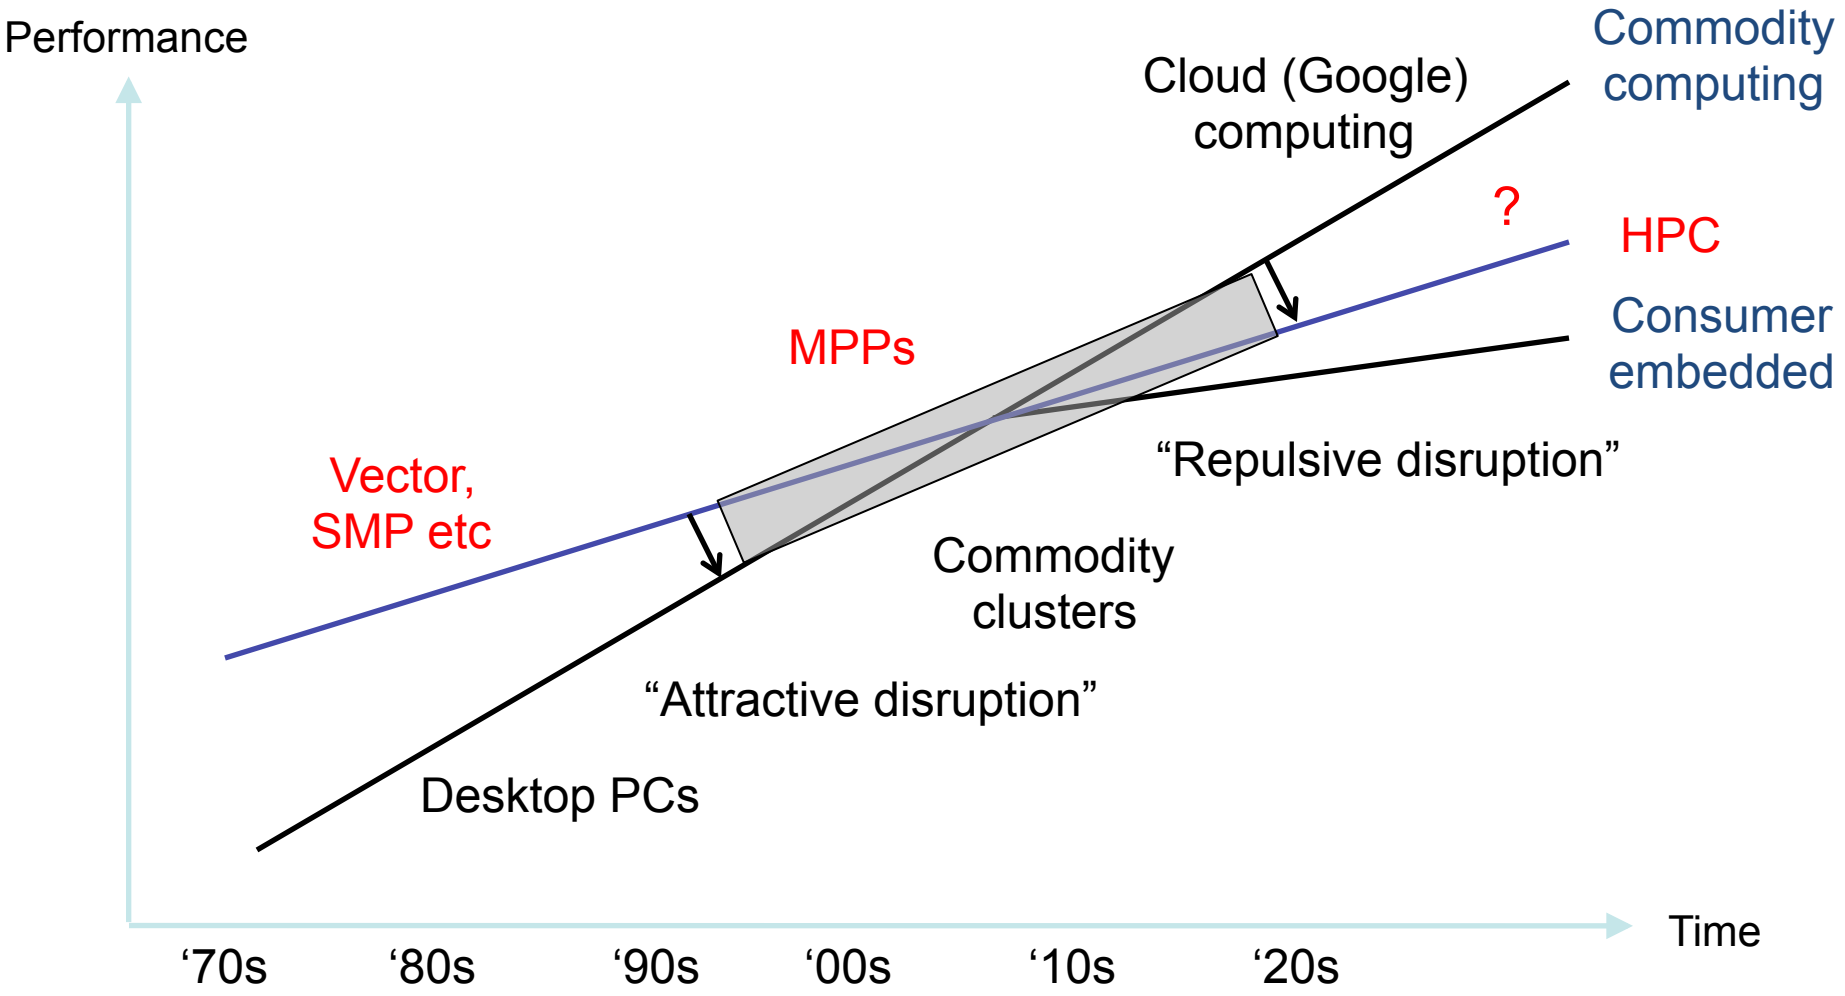

University of BRISTOL

Heterogeneity Changes

University of BRISTOL

Heterogeneity

Causes:

- Multiple types of core
- Interconnect
- Memory type, capacity, …
- Software (OS, middleware, tools, …)

Heterogeneous Systems

Intel MIC

AMD Llano Fusion APUs

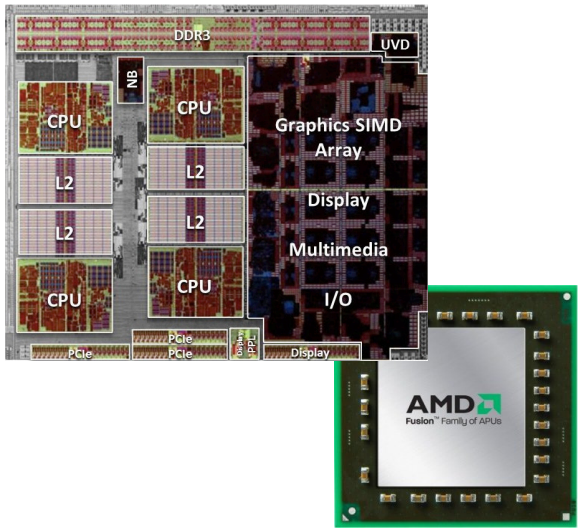

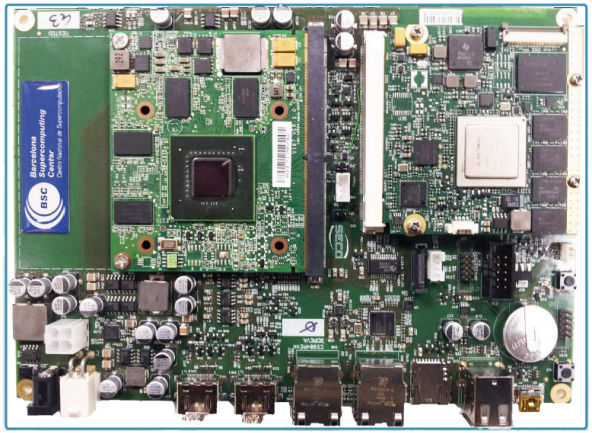

FP7 Mont Blanc ARM + GPU

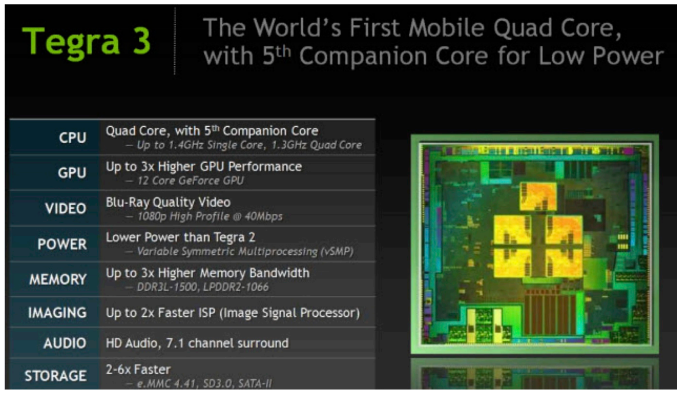

NVIDIA Tegra, Project Denver

University of BRISTOL

Implications

- New programming languages, models, …
- Dynamically adaptive software
  - Discover resources at run-time
- Auto-tuning
- Application frameworks

Scale Changes

University of BRISTOL

Scale

From hundreds to tens of thousands of cores on-chip

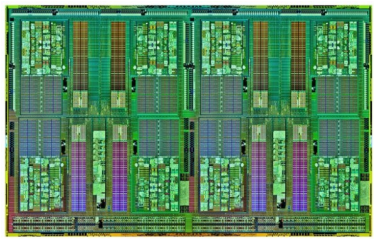

AMD Interlagos
16 cores

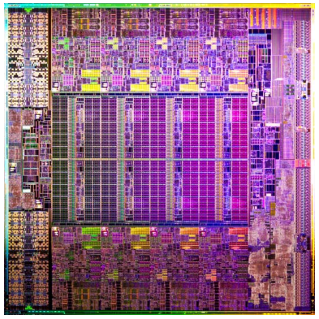

Intel Sandy Bridge E5-2600
8 cores

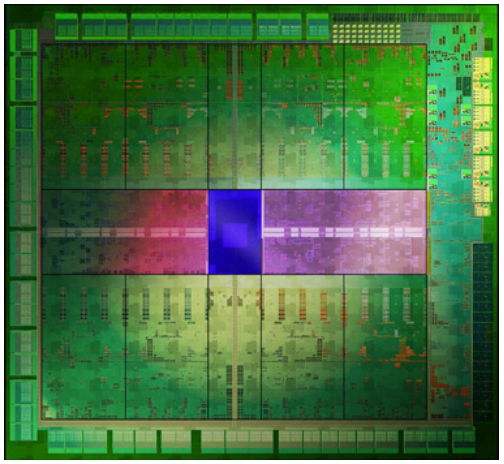

NVIDIA Kepler GTX680
1,536 cores

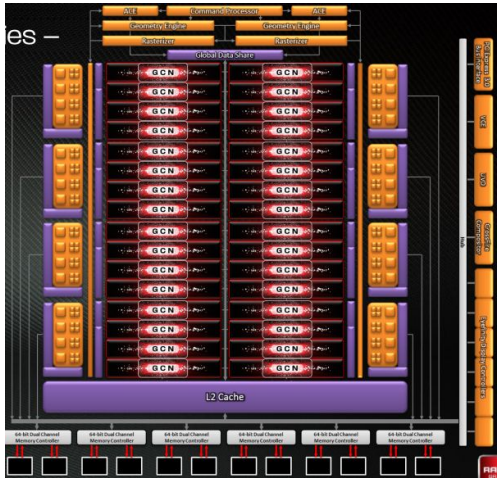

AMD GCN Radeon 7970
2,048 cores

University of BRISTOL

Implications

- ***Massively*** parallel algorithms
  - Billion way parallelism
- Greater reliance on weak scaling?
  - Ruthlessly combat Amdahl's Law with Gustafson's Law
- Islands of performant cache coherence
  - More extreme NUMA

Conclusions

- Five major hardware trends will affect exascale software design:
  1. Changes to memory hierarchies
  2. The impact of fault tolerance
  3. Focus on energy efficiency
  4. Heterogeneity
  5. Scale

- Much of this is already predictable, but not all!
- Lots of work left to do!