Trends in Heterogeneous Systems Architectures (and how they'll affect parallel programming models)

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, 7 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

Heterogeneous Systems Architectures

University of BRISTOL

Moore's Law today

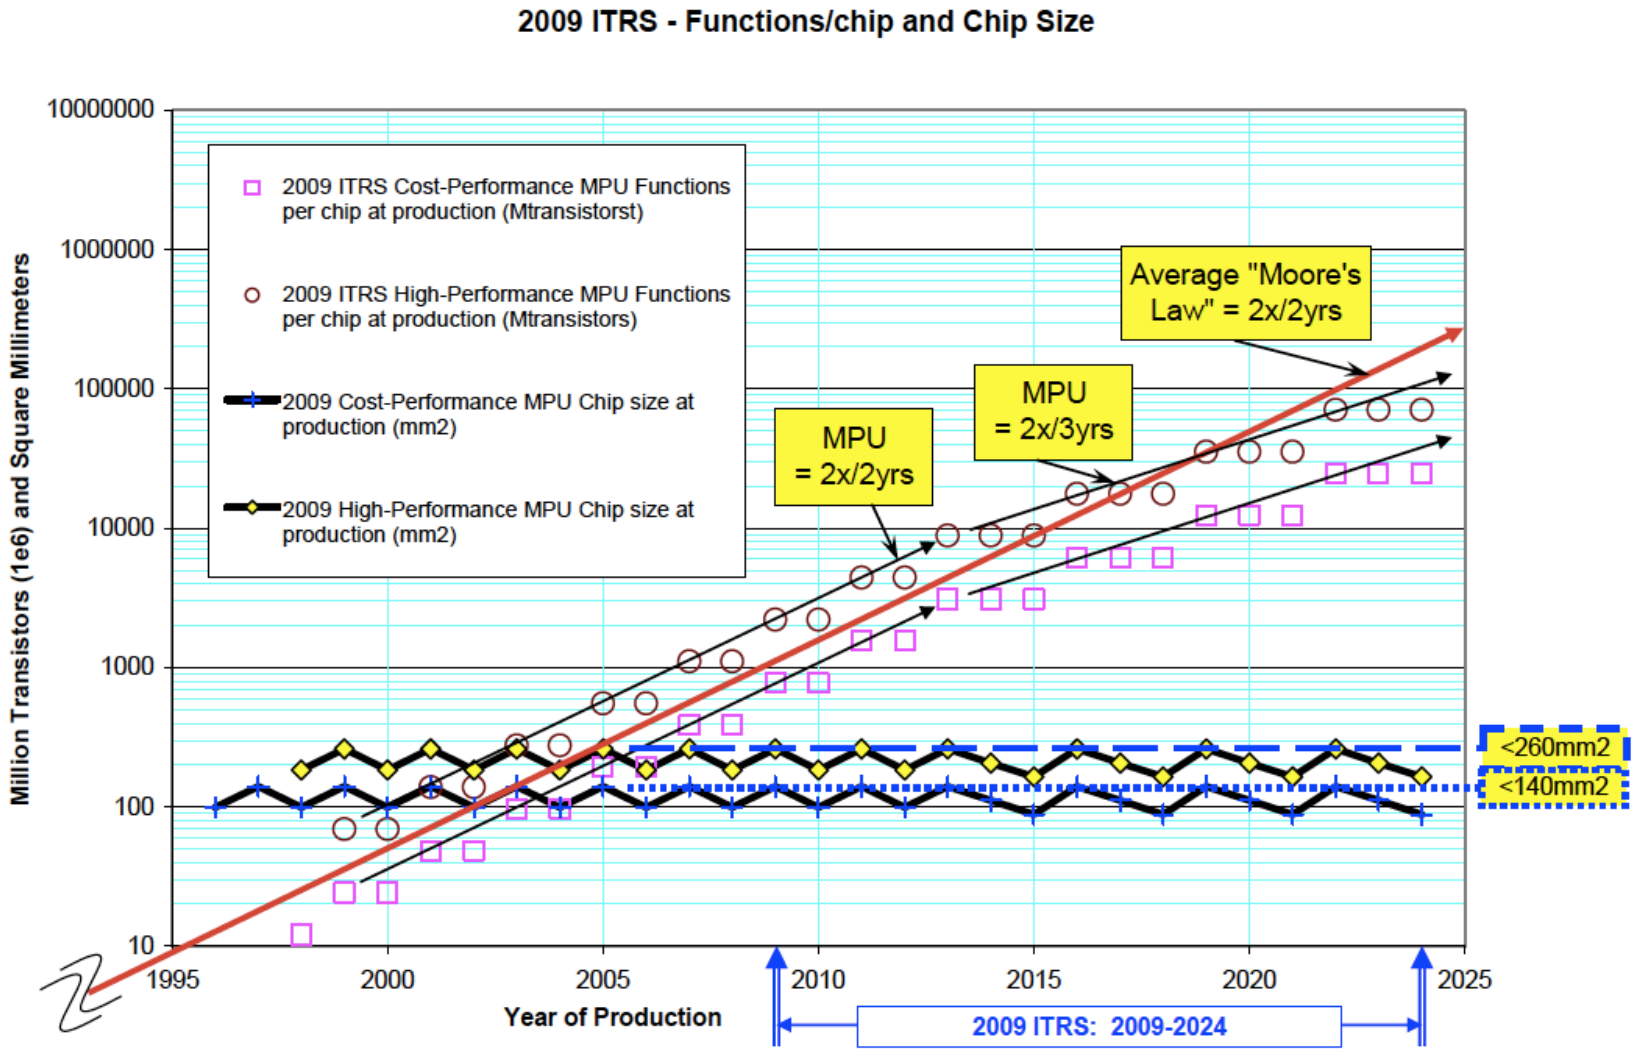

University of BRISTOL

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

Moore's Law today

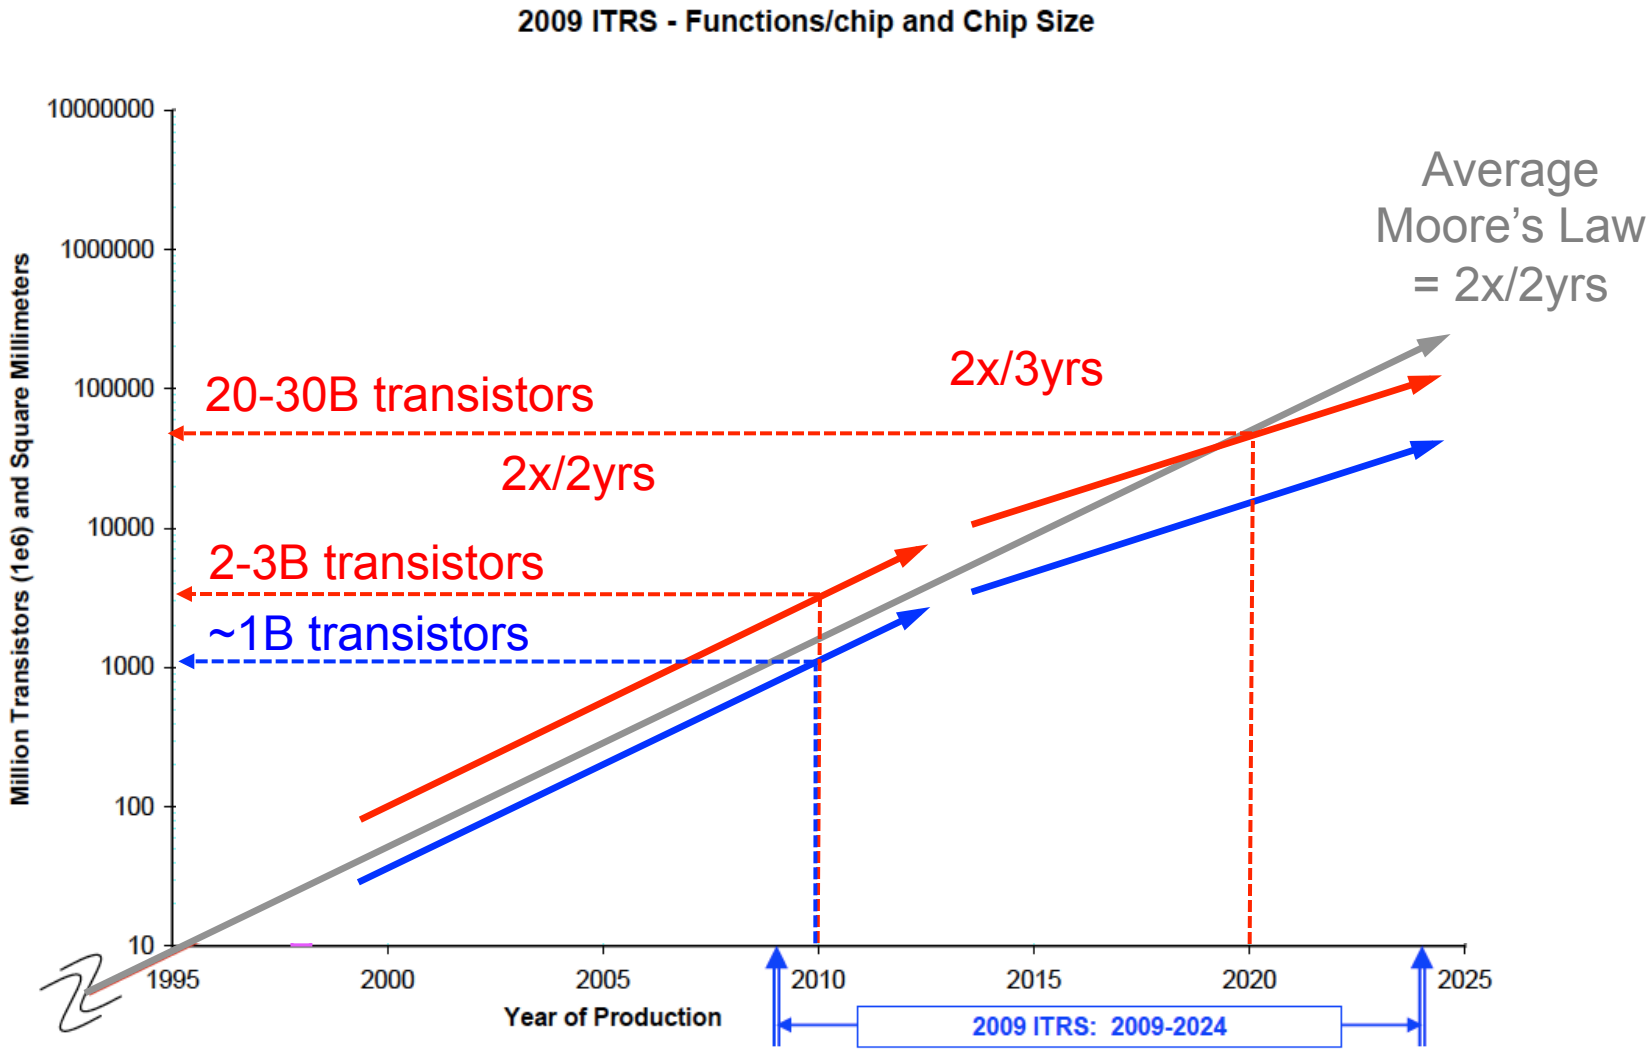

University of BRISTOL

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

Important technology trends

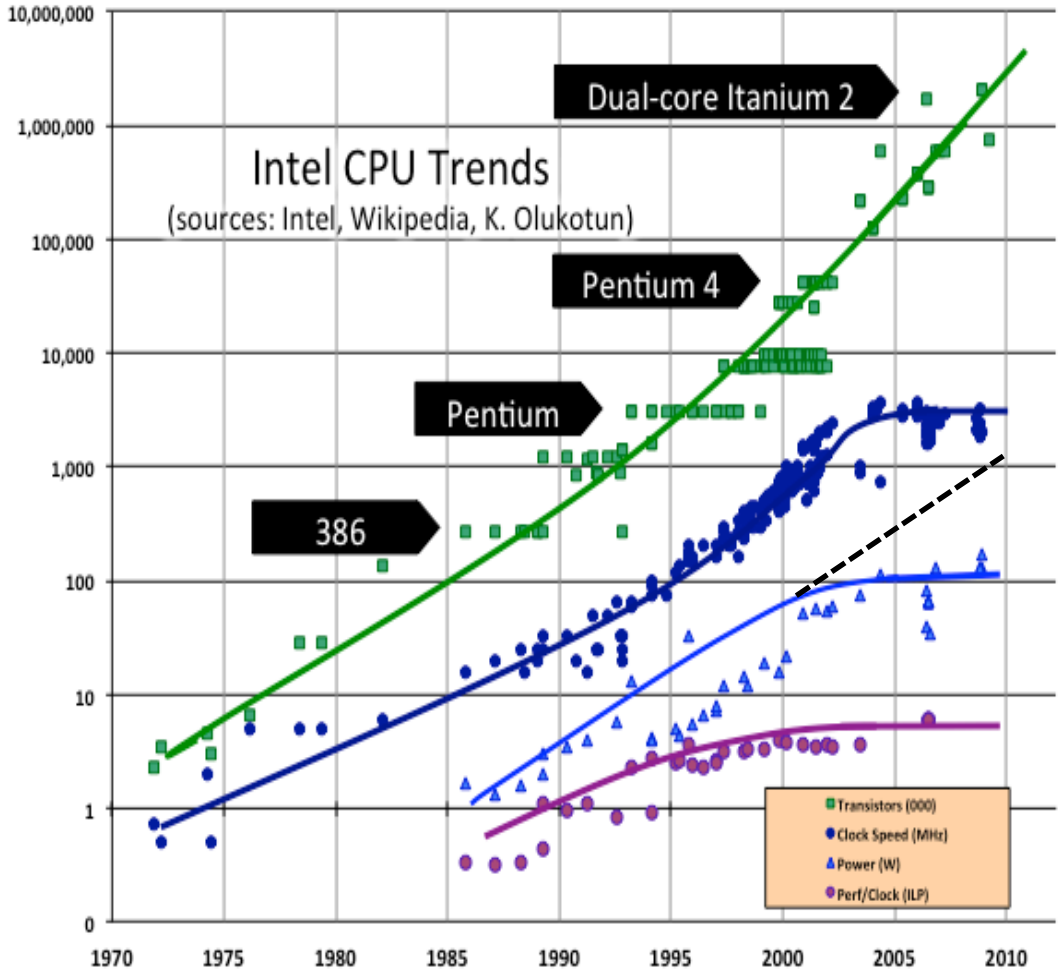

The real Moore's Law

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

Herb Sutter's new outlook

http://herbsutter.com/welcome-to-the-jungle/

In the twilight of Moore's Law, the transitions to multicore processors, GPU computing, and HaaS cloud computing are not separate trends, but aspects of a single trend – mainstream computers from desktops to 'smartphones' are being permanently transformed into heterogeneous supercomputer clusters. Henceforth, a single compute-intensive application will **need to harness different kinds of cores, in immense numbers**, to get its job done.

The free lunch is over. Now welcome to the hardware jungle.

Major hardware trends

University of BRISTOL

The five major hardware trends that will affect Exascale software

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

(For a discussion of 2-5 see "Major hardware trends affecting Exascale developments and their potential impact on software", PRACE- 1IP Work Package 9 Future Technologies Workshop, Daresbury, UK, Apr 11th 2012, http://www.cs.bris.ac.uk/~simonm/publications/sms_prace_exascale_2012.pdf)

Causes of heterogeneity

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

Heterogeneous Systems

Intel MIC

AMD Llano Fusion APUs

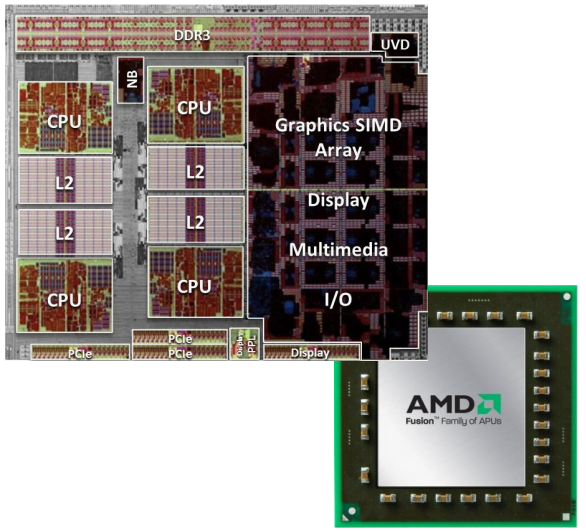

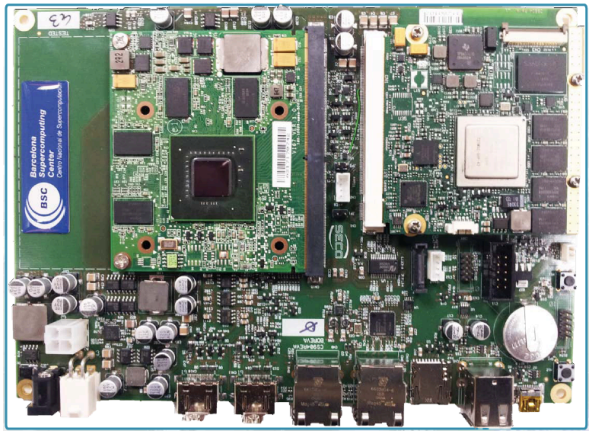

FP7 Mont Blanc ARM + GPU

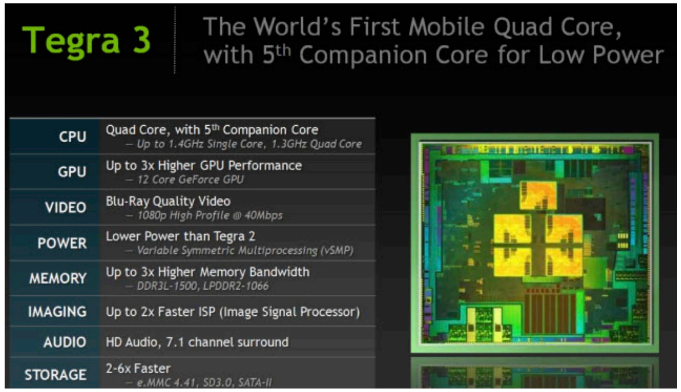

NVIDIA Tegra, Project Denver

University of BRISTOL

Heterogeneity is mainstream

Quad-core ARM Cortex A9 CPU

Quad-core SGX543MP4+ Imagination GPU

You might have one in your pocket

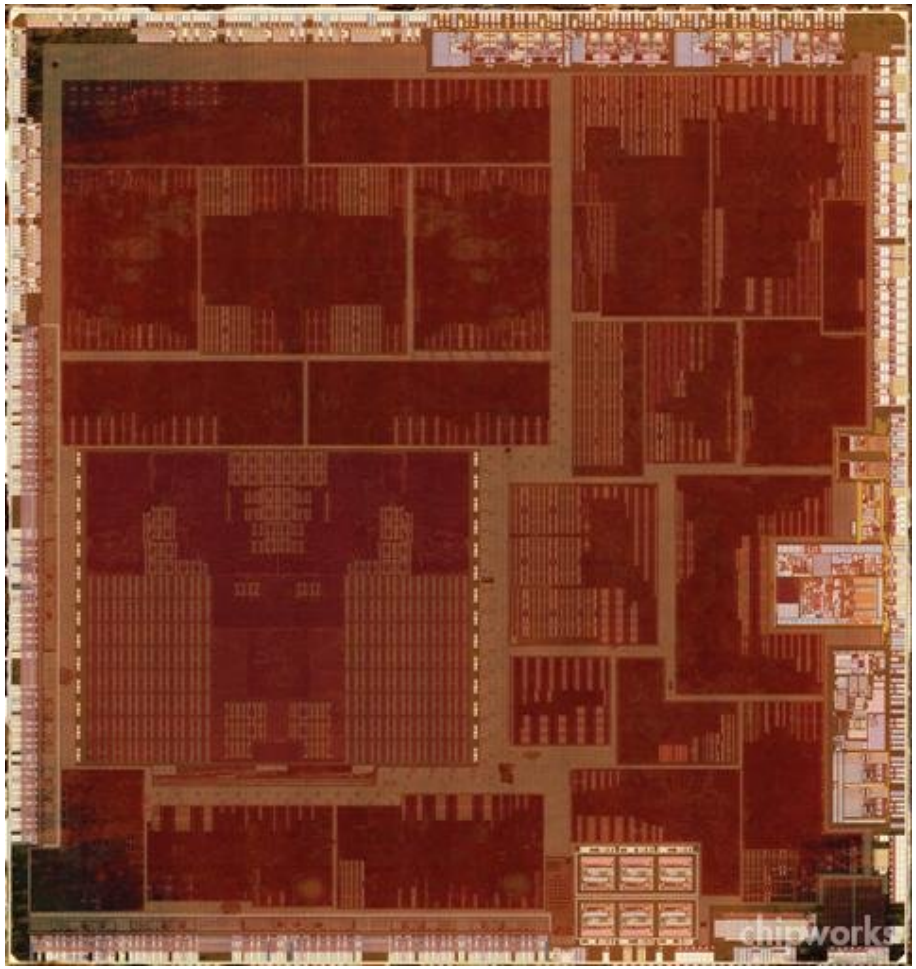

University of BRISTOL

Hot off the press…

INTRODUCING TRINITY: SECOND GENERATION A-SERIES APU

Premium discrete-level graphics and gaming performance

Up to 726 GFLOPs compute

2X the performance/Watt of Llano

56% improvement in graphics performance

Up to 12 hours battery life

All in highly mobile, low-power form factors

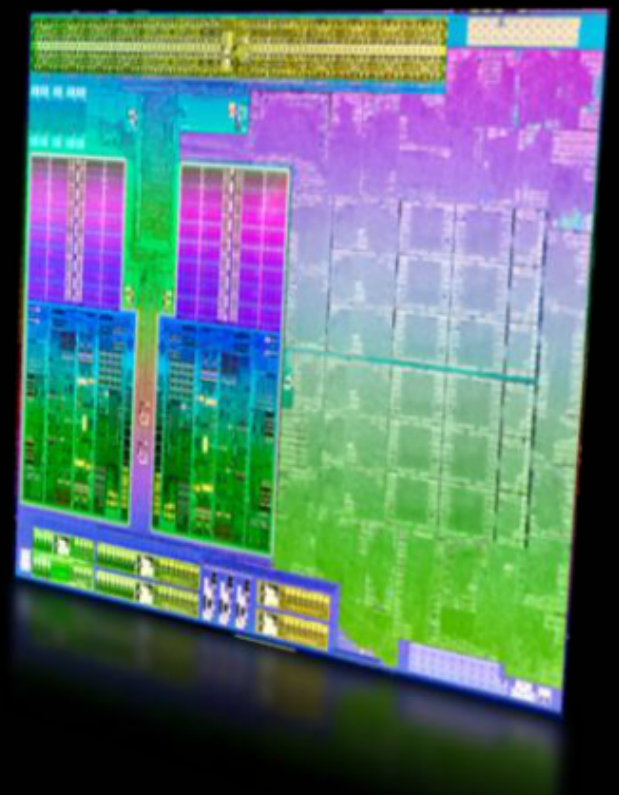

12 | The Promise of Parallel: Today's State of Heterogeneous Computing | June 12, 2012

AMD Fusion 12 DEVELOPER SUMMIT

University of BRISTOL

Implications

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

OpenCL
(Open Computing Language)

University of BRISTOL

OpenCL

- Open standard for portable, parallel programming of heterogeneous systems
- Lets programmers write a single **portable** program that uses **all** resources in the heterogeneous platform

A modern system includes:

- One or more CPUs
- One or more GPUs
- DSP processors
- …other devices?

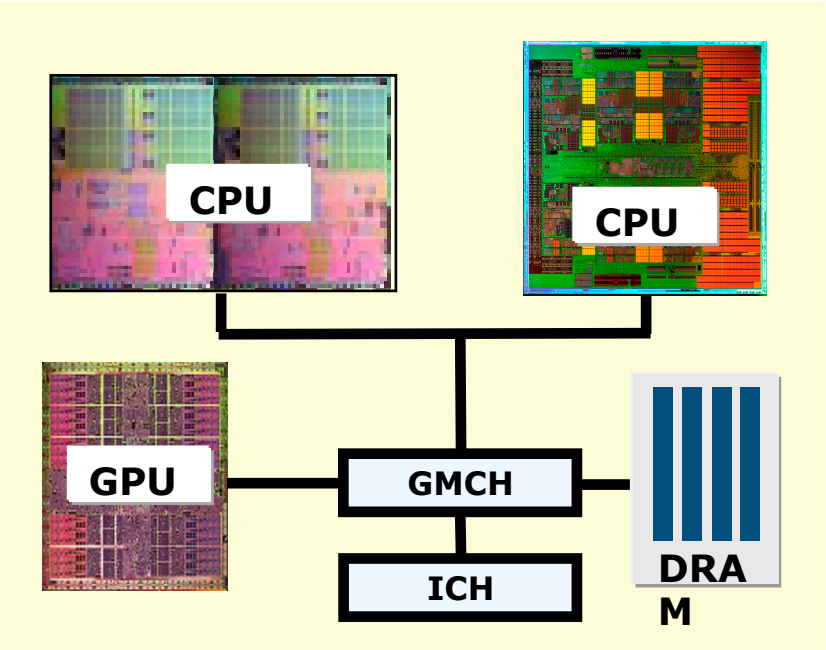

GMCH = graphics memory control hub

ICH = Input/output control hub

OpenCL platform model

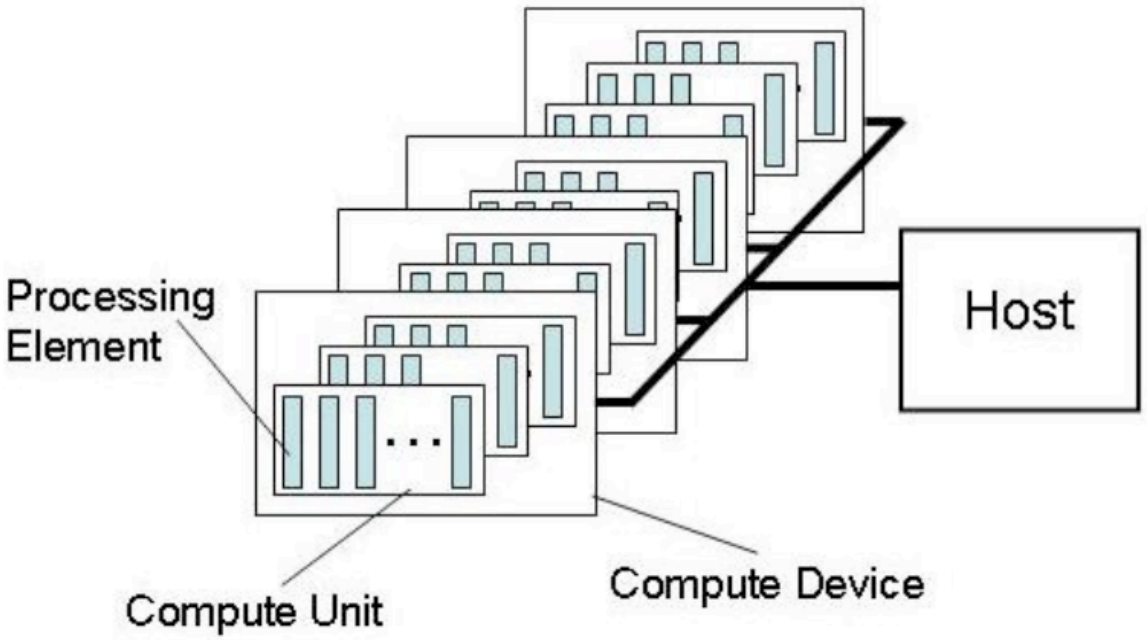

- One Host + one or more Compute Devices
  - Each Compute Device is composed of one or more Compute Units
  - Each Compute Unit is further divided into one or more Processing Elements

The BIG idea behind OpenCL

- Replace loops with functions (a kernel) executing at each point in a problem domain (index space).
- E.g., process a 1024 x 1024 image with one kernel invocation per pixel or 1024 x 1024 = 1,048,576 kernel executions

Traditional loops

```
void
trad_mul(const int n,
         const float *a,
         const float *b,
         float *c) {
  int i;
  for (i=0; i<n; i++)
    c[i] = a[i] * b[i];
}
```

Data parallel OpenCL

```
kernel void
dp_mul(global const float *a,
       global const float *b,
       global float *c) {
  int id = get_global_id(0);

  c[id] = a[id] * b[id];

} // execute over "n" work-items
```

OpenCL memory model

- Private Memory
  - Per Work-Item
- Local Memory
  - Shared within a Work-Group
- Global / Constant Memories
  - Visible to all Work-Groups
- Host Memory
  - On the CPU

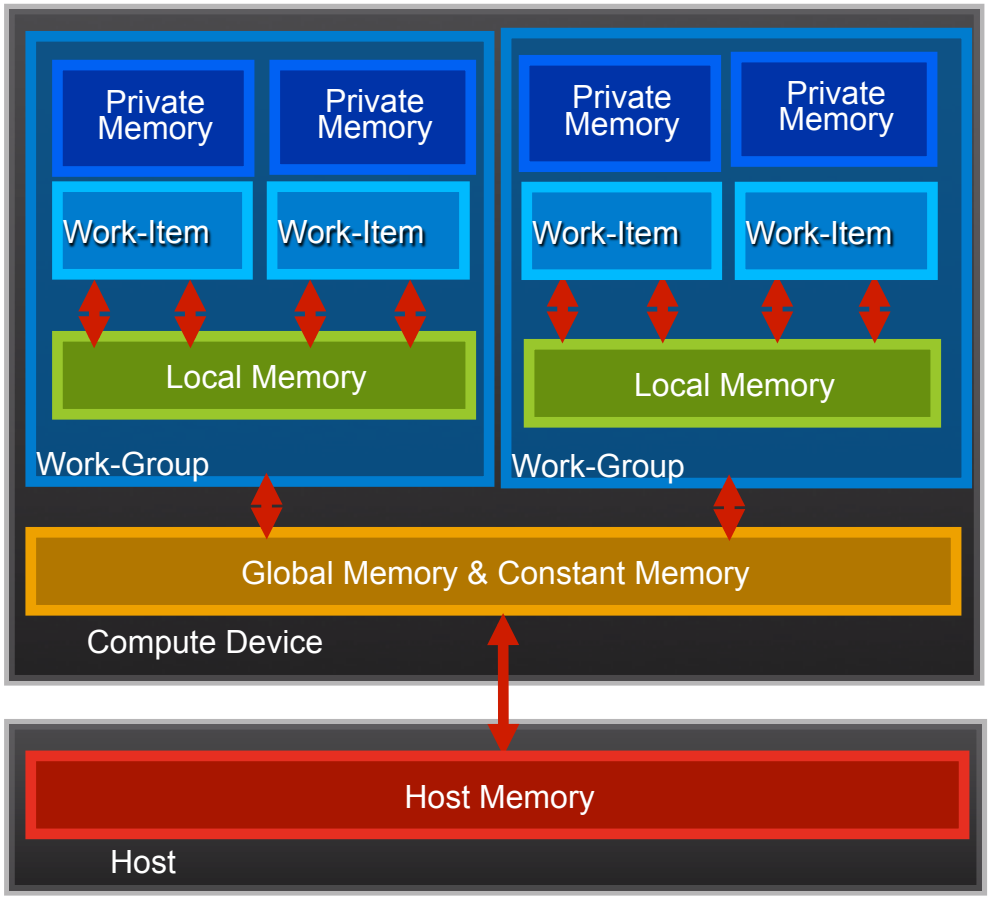

Memory management is explicit
You must move data from host → global → local *and* back

University of BRISTOL

Issues with OpenCL*

- It does not **compose**
  - Disjoint memory address spaces (local/global)
  - Barriers

- It provides no resource management
  - Kernels are a statically allocated resource

* And most other parallel programming languages

Composability

Need to support many algorithms, even within a single application

Task farms, pipeline, data parallelism, …

See similar composability challenges with OpenMP and parallel libraries, for example

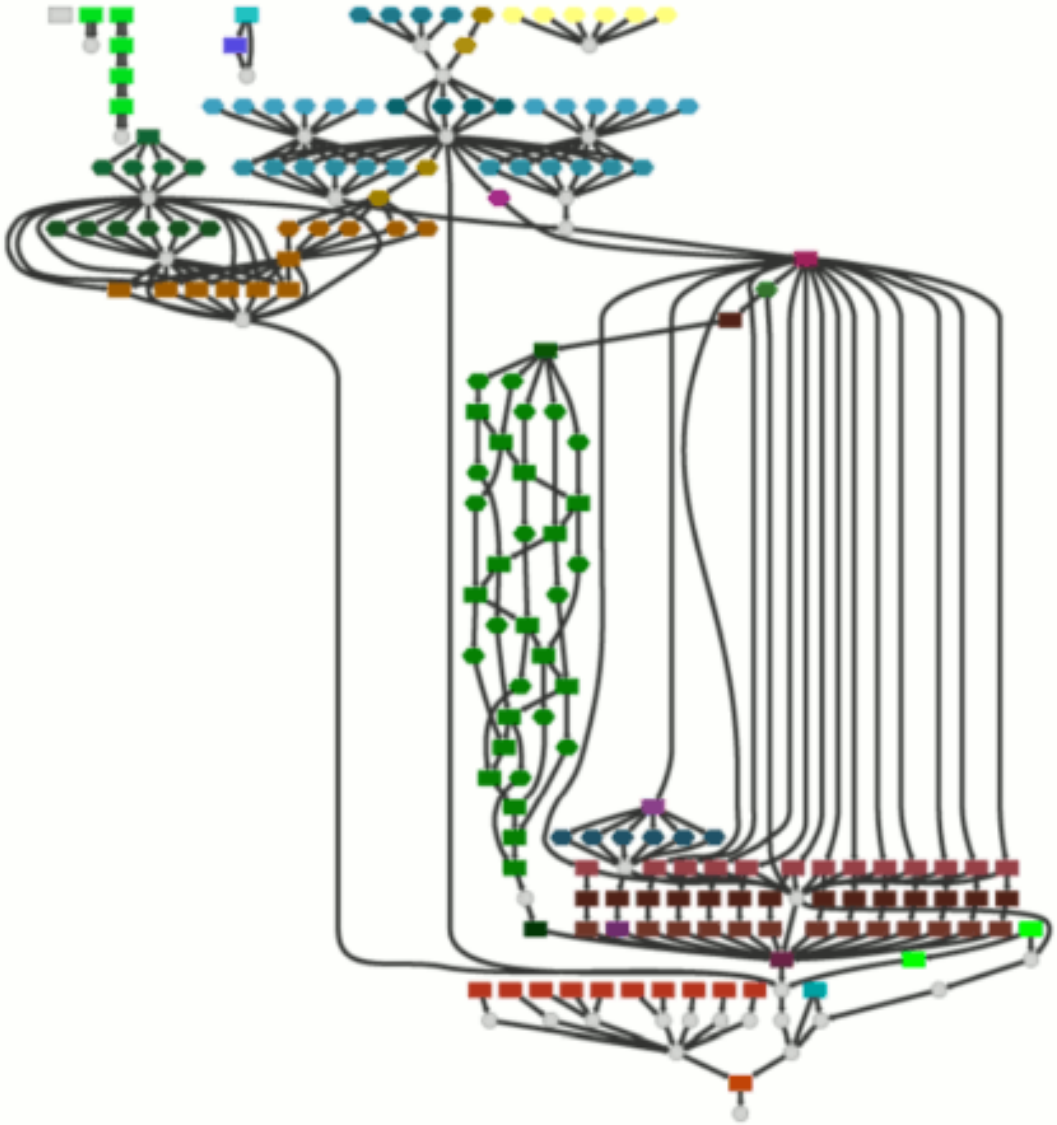

Heterogeneous System Architecture (HSA)

University of BRISTOL

HSA overview

- The HSA Foundation launched this week and already includes AMD, ARM, Imagination Technology and Texas Instruments

- HSA is a new, ***open*** architecture specification
  - HSAIL virtual ISA
  - HSA memory model
  - HSA dispatch

- Provides an optimised platform architecture for heterogeneous programming models such as OpenCL, C++AMP, et al

Announced 12/6/2012

THE HSA FOUNDATION: ACTIVITIES

- Nonprofit, open standardization body for HSA platforms that will own the development and evangelization of the architecture going forward
- Make heterogeneous programming easy and a first-class pervasive complement to CPU computing
- Continue to increase the power efficiency of HSA, keeping it the platform of choice from smartphones to the cloud
- Bring to market strong development solutions (tools, libraries, OS runtimes) to drive innovative advanced content and applications
- Foster growth of heterogeneous computing talent through HSA developer training and academic programs to drive both learning and innovation

HSA overview

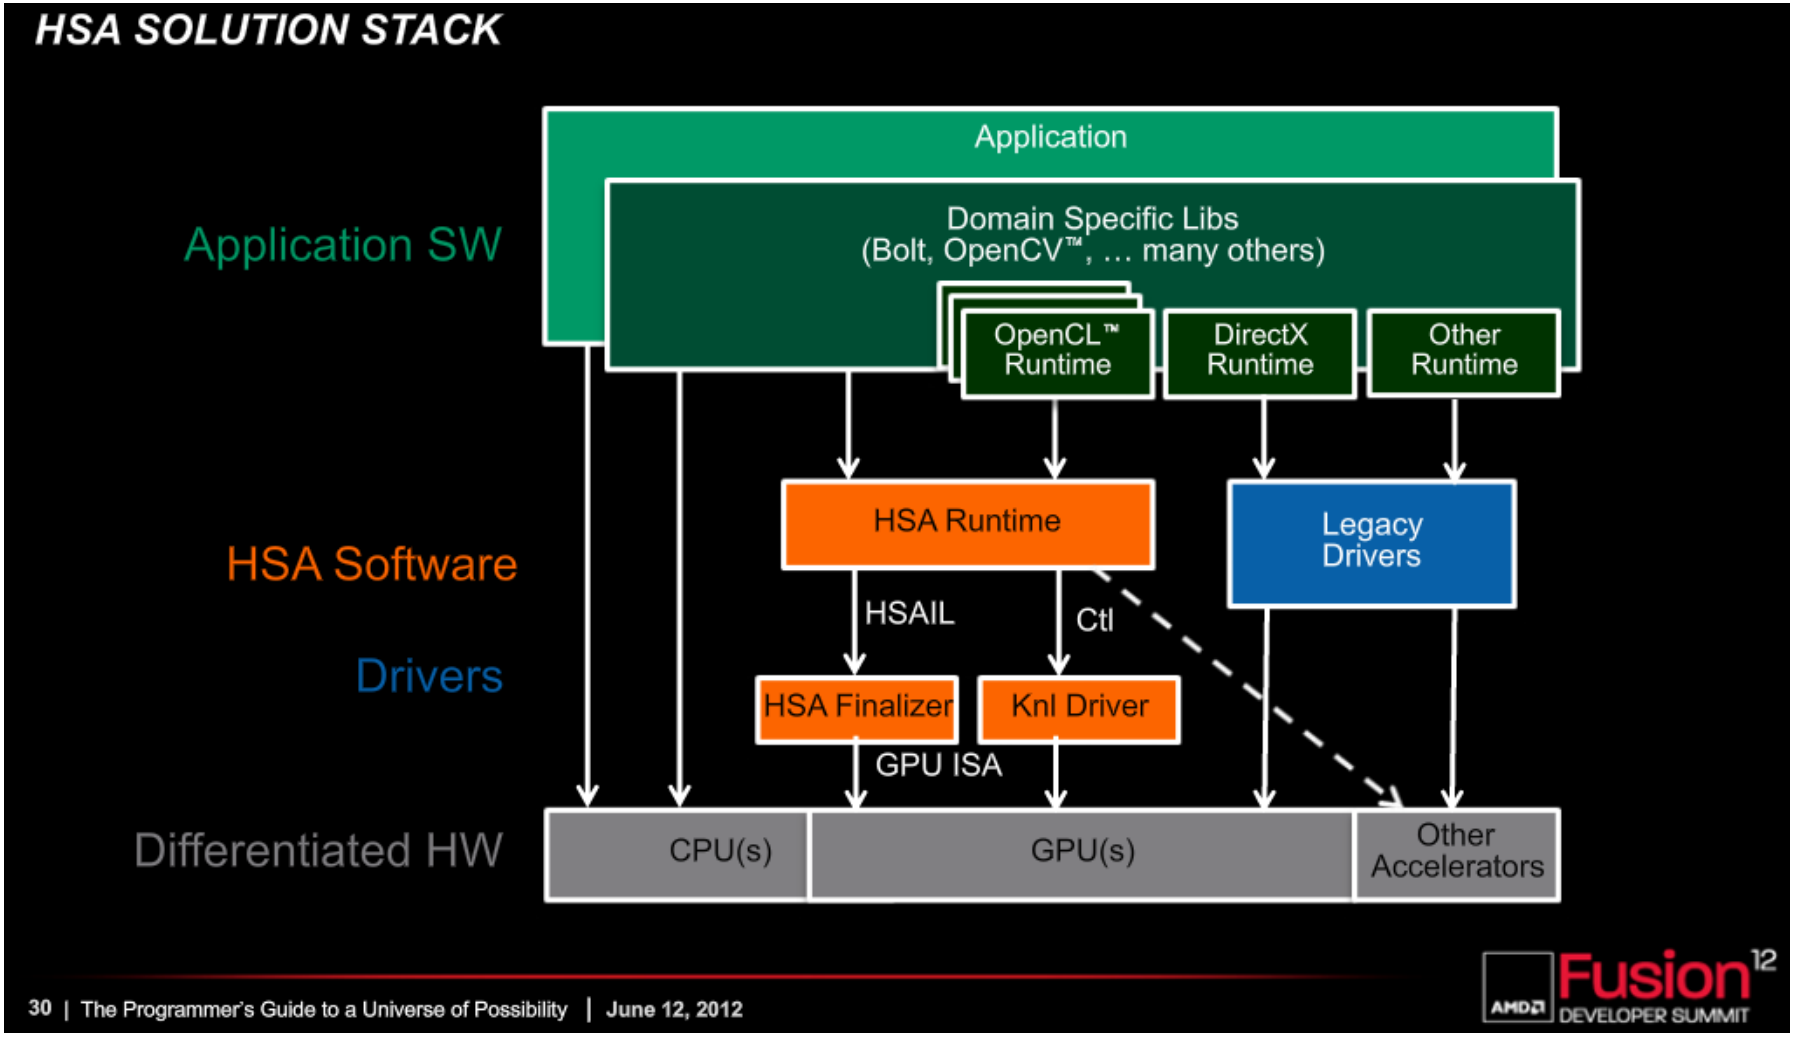

University of BRISTOL

HSA features: simplifying programming

- Much finer grained integration of CPU and GPU cores in silicon
- Unified address space for all cores
- Will support GPU context switching, preemption
- PGAS-style distributed arrays
  - Memory hierarchy abstraction to address function composition
- First class barrier objects
  - Aids composability

HSA Intermediate Layer (HSAIL)

- Virtual ISA for parallel programs
- Similar idea to LLVM IR - a good target for compilers
- *Finalised* to specific ISA by a JIT compiler
- Features:
  - Explicitly parallel
  - Support for exceptions, virtual functions and other high-level features
  - Syscall methods (I/O, printf etc.)
  - Debugging support

HSA memory model

- Compatible with C++11, Java and .NET memory models
- Relaxed consistency
- Designed to support both managed language (such as Java) and unmanaged languages (such as C)
- Will make it much easier to develop 3rd party compilers to target a wide range of heterogeneous products
  - Fortran, C++, C++AMP et al

HSA dispatch

- HSA designed to enable heterogeneous task queuing
  - A work queue per core
  - Distribution of work into queues
  - Load balancing by work stealing

- Any core can schedule work for any other, including itself
- Significant reduction on overhead of scheduling work for a core

Today's command and dispatch flow

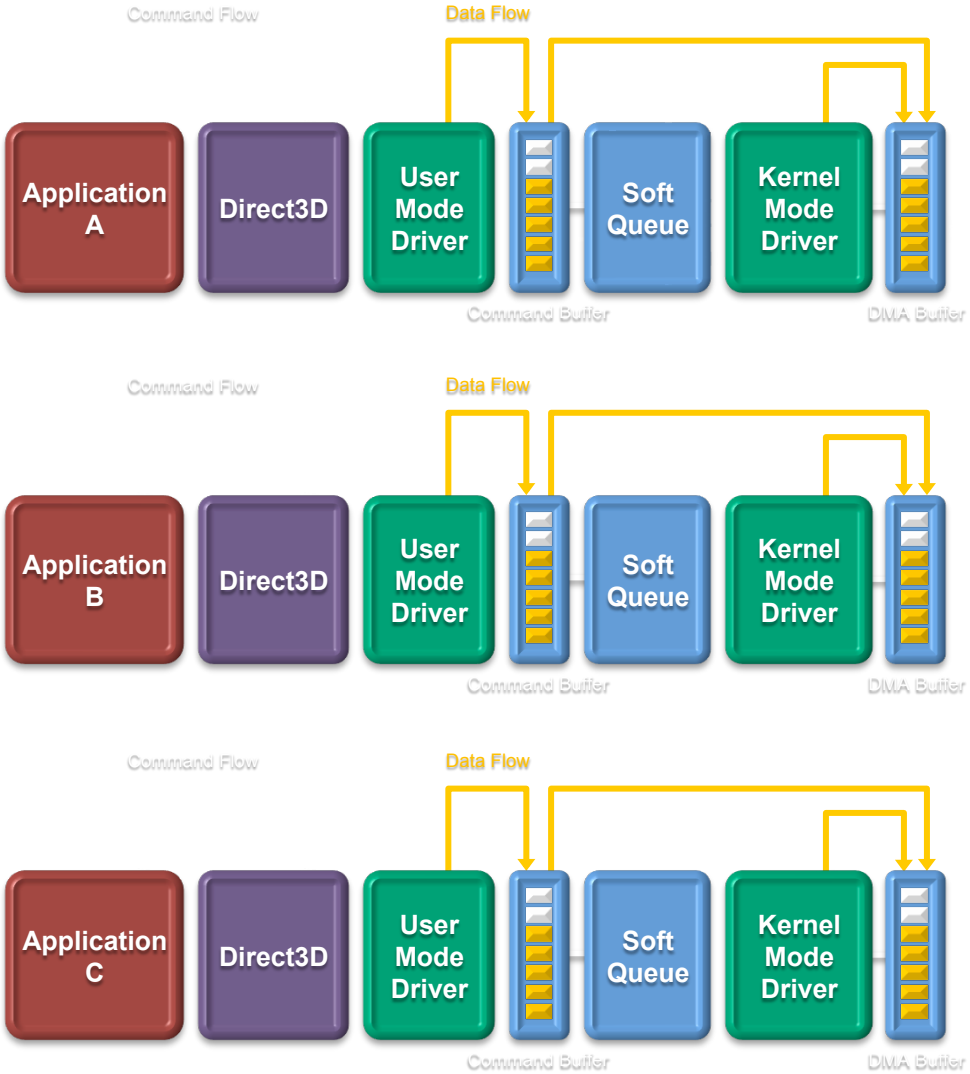

University of BRISTOL

HSA roadmap from AMD

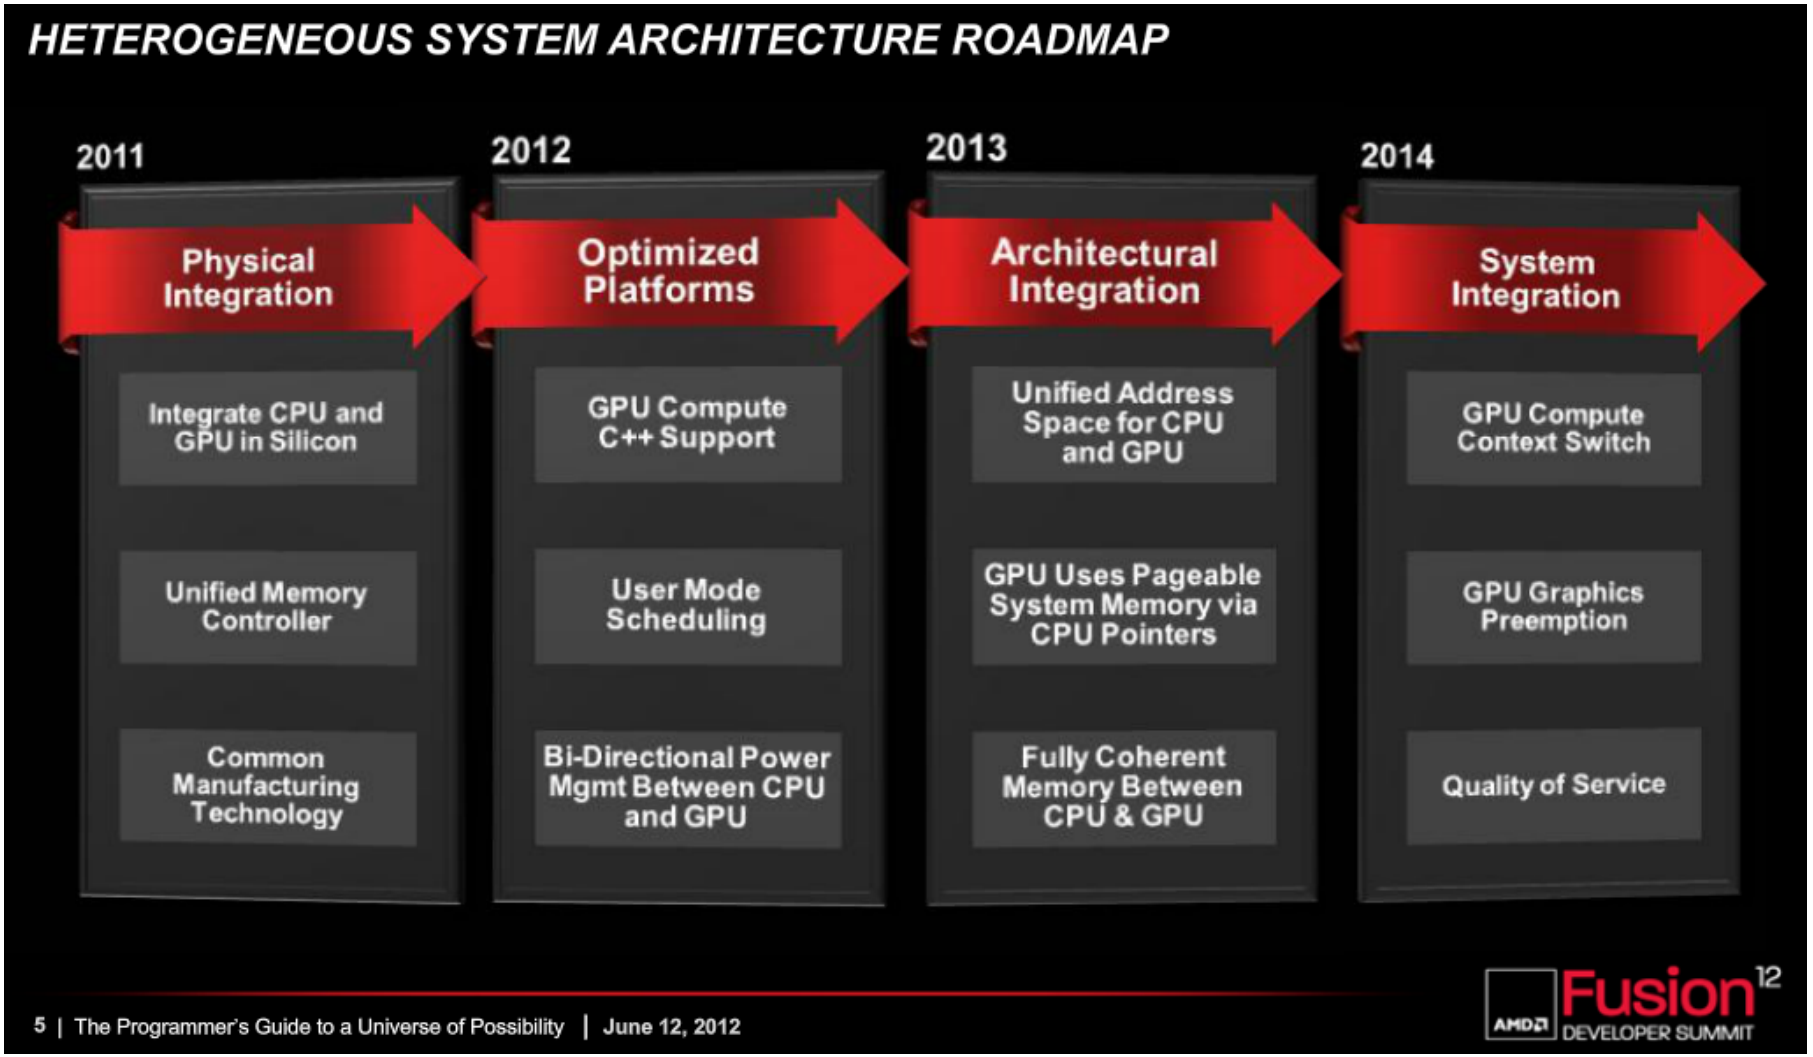

HSA tools released as open source

AMD'S OPEN SOURCE COMMITMENT TO HSA

- We will open source our linux execution and compilation stack
  - Jump start the ecosystem
  - Allow a single shared implementation where appropriate
  - Enable university research in all areas

Component Name	AMD Specific	Rationale
HSA Bolt Library	No	Enable understanding and debug
OpenCL HSAIL Code Generator	No	Enable research
LLVM Contributions	No	Industry and academic collaboration
HSA Assembler	No	Enable understanding and debug
HSA Runtime	No	Standardize on a single runtime
HSA Finalizer	Yes	Enable research and debug
HSA Kernel Driver	Yes	For inclusion in linux distros

Conclusions

- Heterogeneity is an increasingly important trend
- The market is finally starting to create and adopt the necessary open standards
  - Proprietary models likely to start declining now
  - Don’t get locked into any one vendor!
- Parallel programming models are likely to (re)proliferate
- Exciting times ahead!

www.cs.bris.ac.uk/Research/Micro