

Performance Portability for Embedded GPUs

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





GPUcompute on Embedded GPUs

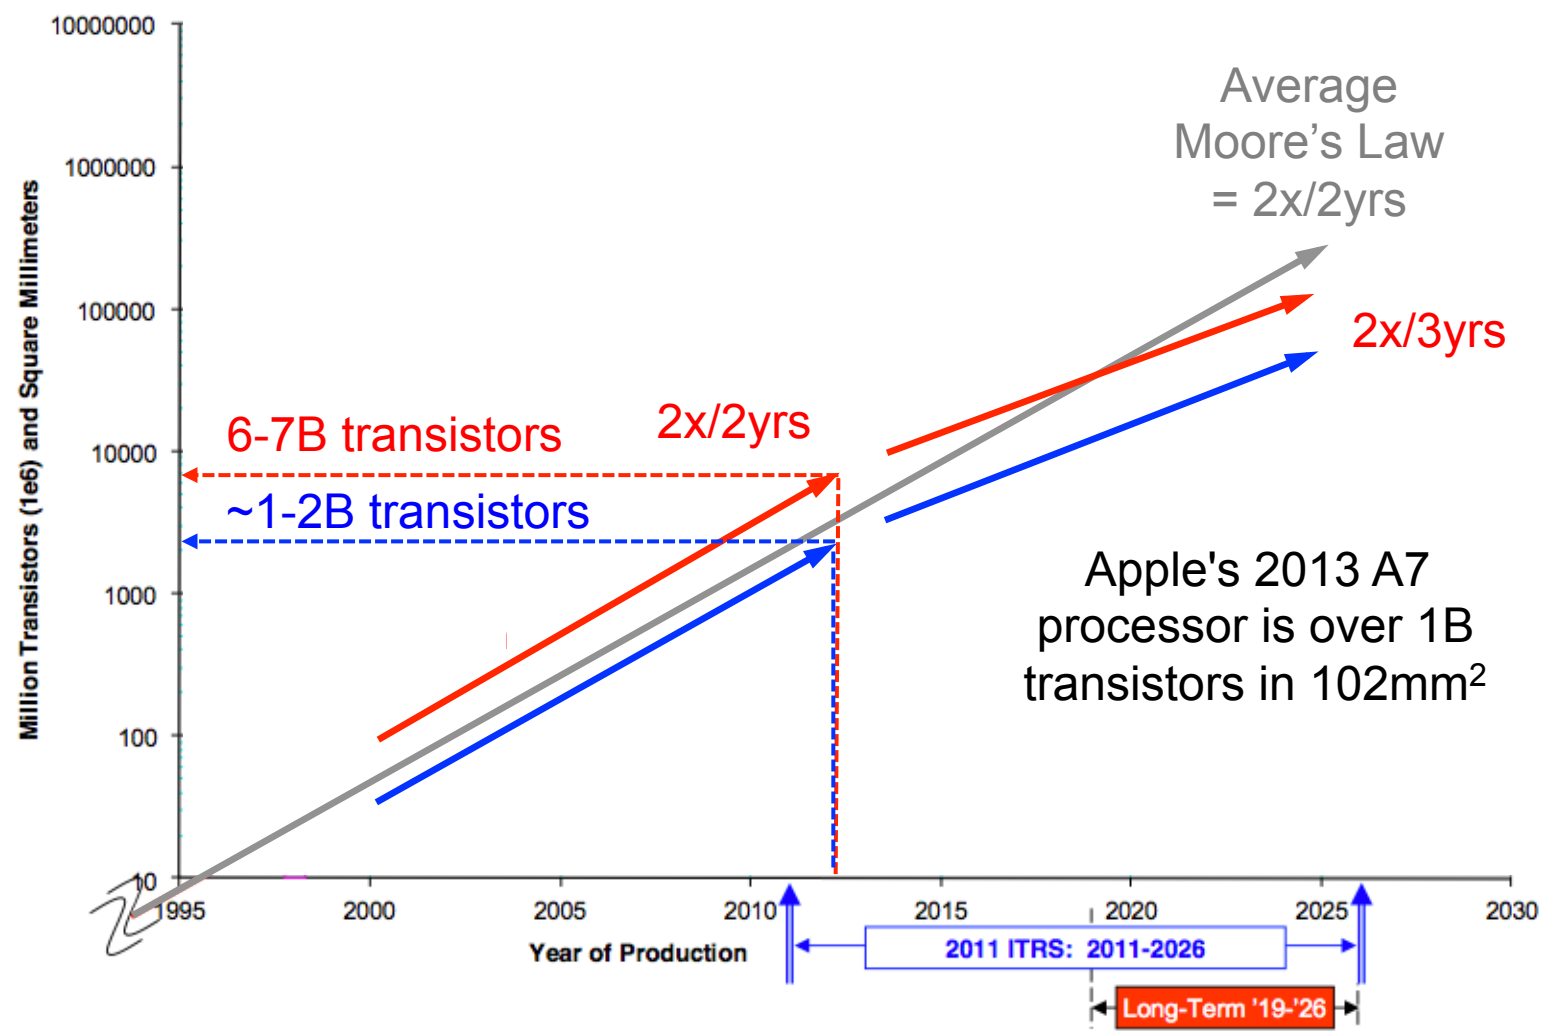
Simon McIntosh-Smith simonm@cs.bris.ac.uk
Head of Microelectronics Research
University of Bristol, UK





Embedded GPUs becoming monsters

2011 ITRS - Functions/chip and Chip Size



Programmable embedded GPUs go mainstream in 2014

- Lots of significant announcements recently (at CES etc)
 - Imagination Series 6XT (Rogue) GPU
 - ARM Mali T760
 - Qualcomm Adreno 420
 - AMD Radeon E6XXX (576 GFLOPS), Kaveri
 - Intel IRIS
 - Vivante GC7000 series
 - Nvidia Tegra K1
 - Other vendors too: Broadcom, ...
- 2014 is the year that OpenCL programmable embedded GPUs will become ubiquitous

Architectural trends

- Integrated CPU/GPU with shared virtual memory and cache coherency
 - SVM is a key new feature of OpenCL 2.0
- Efficiency improvements in overheads for launching tasks/kernels
 - Promotes finer-grained heterogeneous parallel programs
 - Supported by OpenCL 2.0 subgroups and nested (dynamic) parallelism
 - Heterogeneous Systems Architecture (HSA) etc

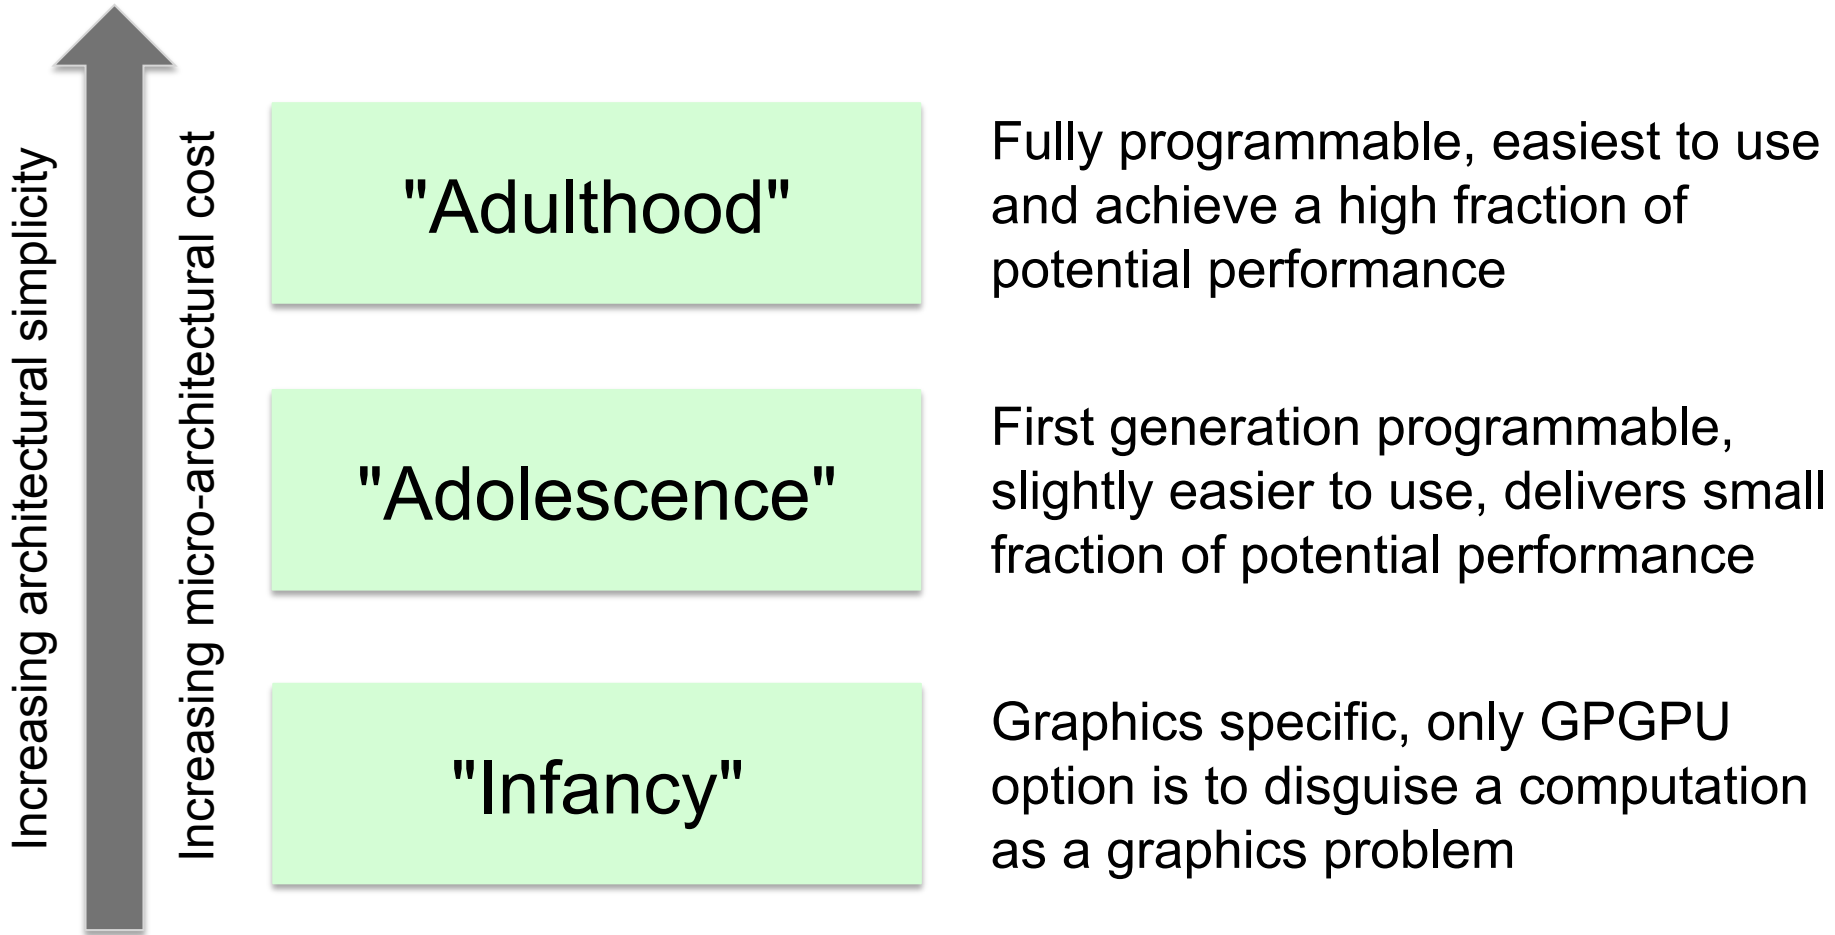
Where are we now?

- Usefully (OpenCL) programmable GPUs just starting to become available
- OpenCL 2.0, SPIR and HSA will enable much greater efficiency in the exploitation of heterogeneous many-core *integrated* systems
- Embedded GPUs having peak performance in the tens to hundreds of GFLOPS (single precision)
- Memory bandwidth of O(10) Gbytes/s
- **But...**

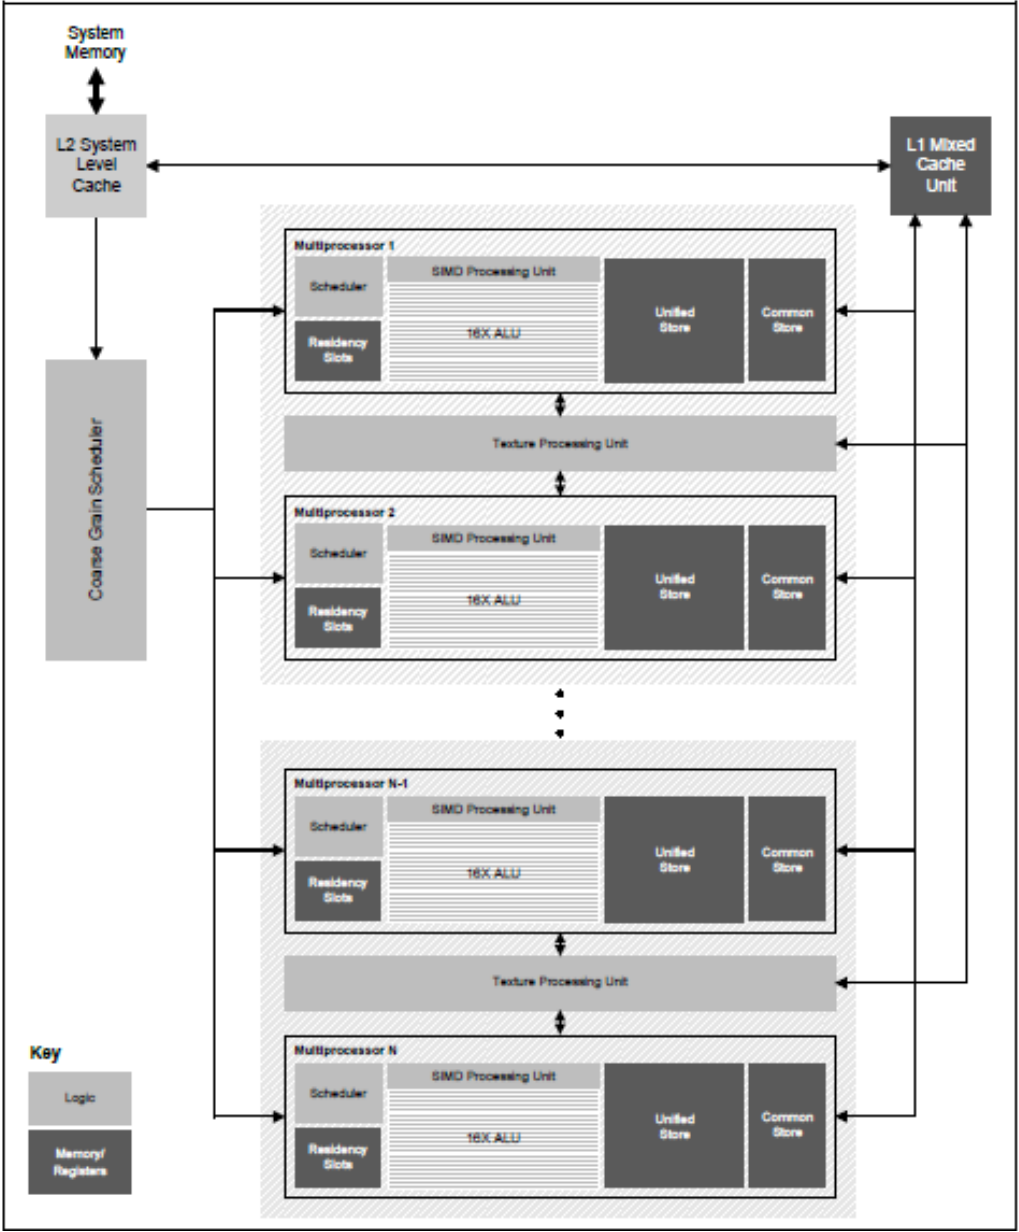
eGPUs still hard to access

- Few eGPUs come with easy to use, familiar programming environments or even operating systems for HPC
 - Most are Android
 - Only a few provide Linux
- Many of their architectures are still in the GPU computing "adolescence" stage

GPU computing evolution



🔥 E.g. Imagination Series 6 "Rogue"



Sources:

<http://anandtech.com/show/5364/powervr-series-6-rogue-gpus-released-to-licensing>

<http://withimagination.imgtec.com/index.php/powervr/powervr-rogue-designing-an-optimal-architecture-for-graphics-and-gpu-compute>

<http://withimagination.imgtec.com/index.php/powervr/building-efficient-multimedia-architectures-consumer-electronics-mobile-computing>

🔥 Dev platforms are hard to find...

ARM Mali-T604 (all Samsung Exynos 5250):

- Arndale development board
- Samsung Chromebook
- Google/Samsung Nexus 10 tablet



Qualcomm Adreno 320 or 330:

- Sony Xperia Z, Xperia ZL, Xperia Tablet Z, Xperia ZR, Xperia Z1 and Xperia Z Ultra smartphones



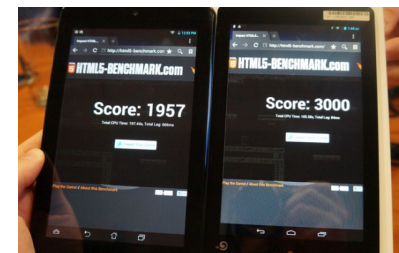
Imagination PowerVR Series 5 (SGX544MP3):

- Hardkernel ODROID-XU development board
- Samsung Galaxy S4 smartphone

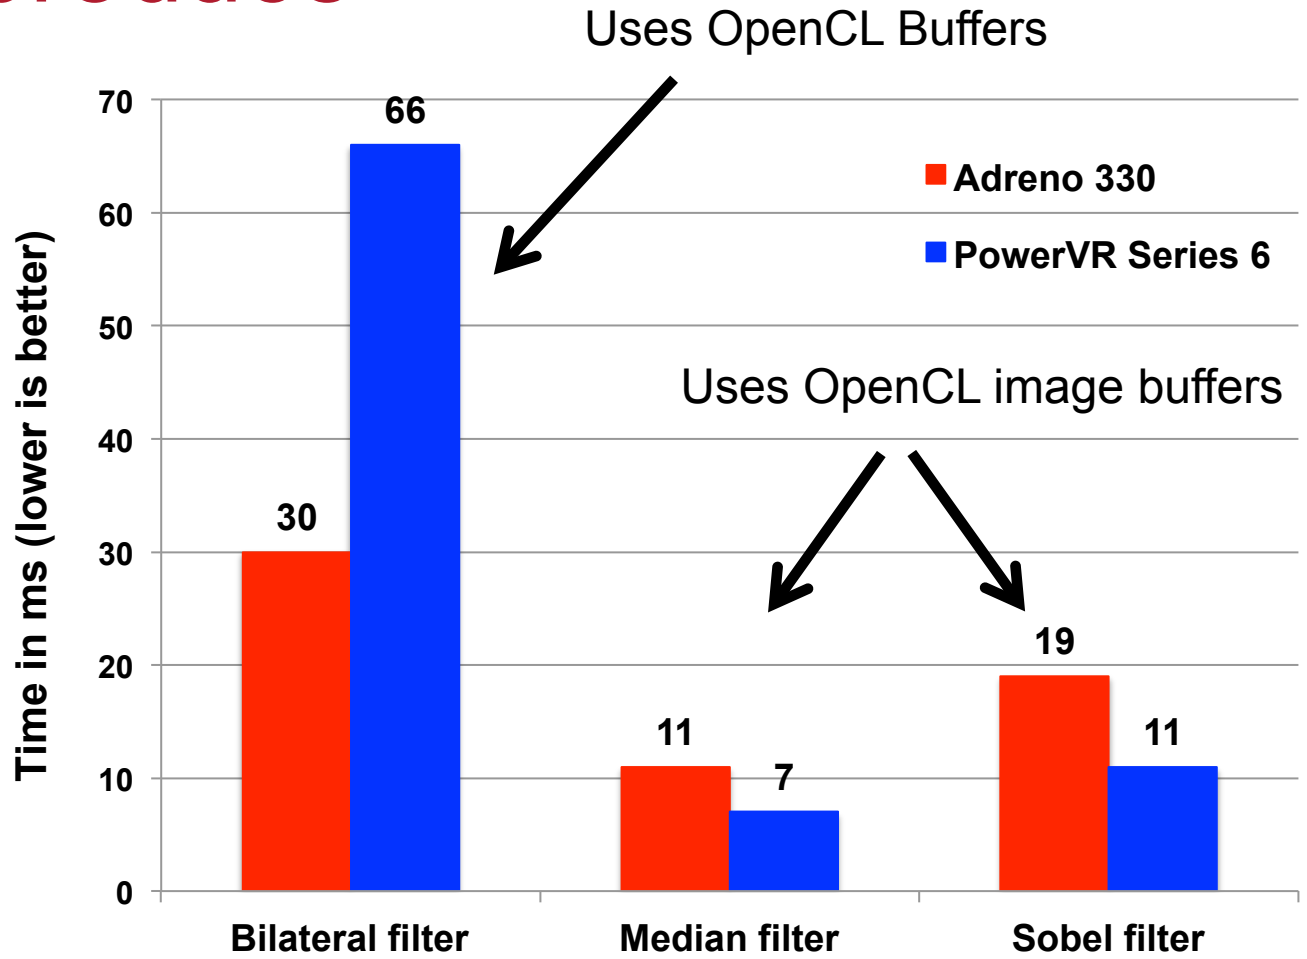


Imagination PowerVR Series 6 ("Rogue"):

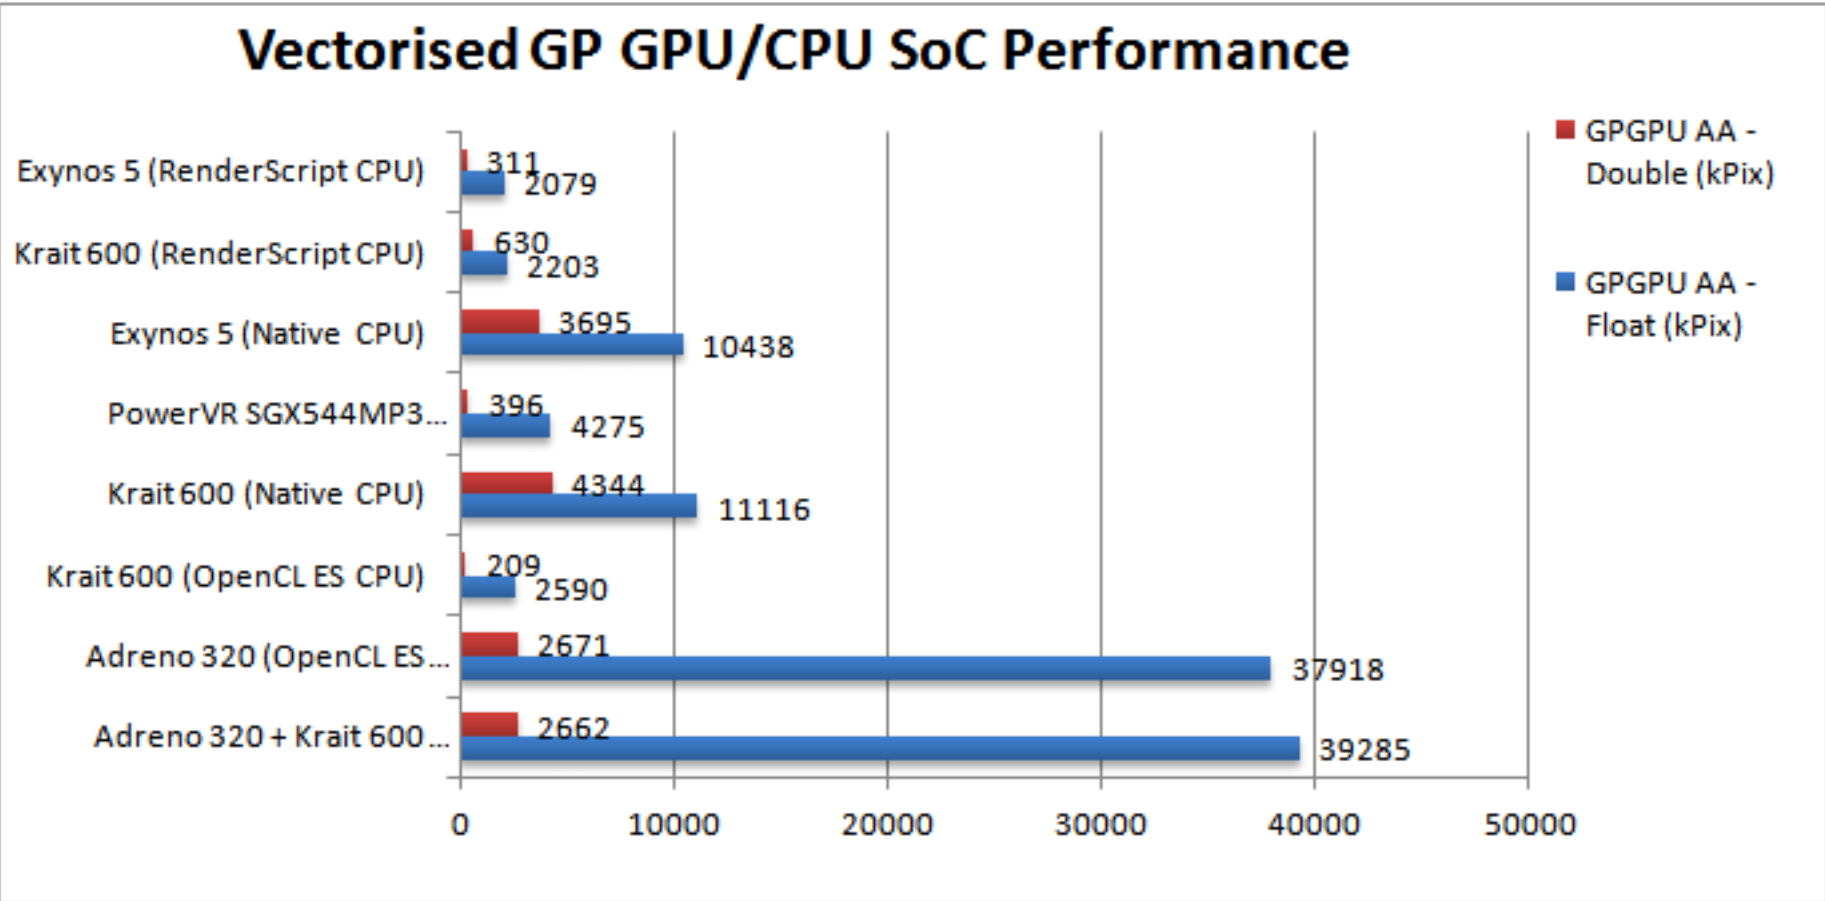
- Articles online claim tablets based on MediaTek's new quad-core MT8135 SoC will be appearing imminently




🔥 Meaningful comparisons still difficult to produce



🔥 OpenCL ES results from SiSoft



🔥 OpenCL ES results from SiSoft



OpenCL ES results from SiSoft

SiSoft's OpenCL ES 2014 benchmark conclusion:


"The good news is that the OpenCL run-time works well for all devices, with fast transfers comparable to SIMD optimised transfers (internal copy). The bad news is that upload/download is diabolically slow and here zero-copy / HSA (Heterogeneous System Architecture) is badly needed to reduce bandwidth pressure."

CompuBench OpenCL results


Only 2 platforms listed:

- Google Nexus 10 (Android 4.2.2)
 - GPU: ARM Mali T604
- Sony Xperia Z
 - GPU: Qualcomm Adreno 320
- Source:
 - <http://compubench.com/result.jsp>
 - Uncheck “Desktop” and “Notebook” on the right

CompuBench OpenCL results

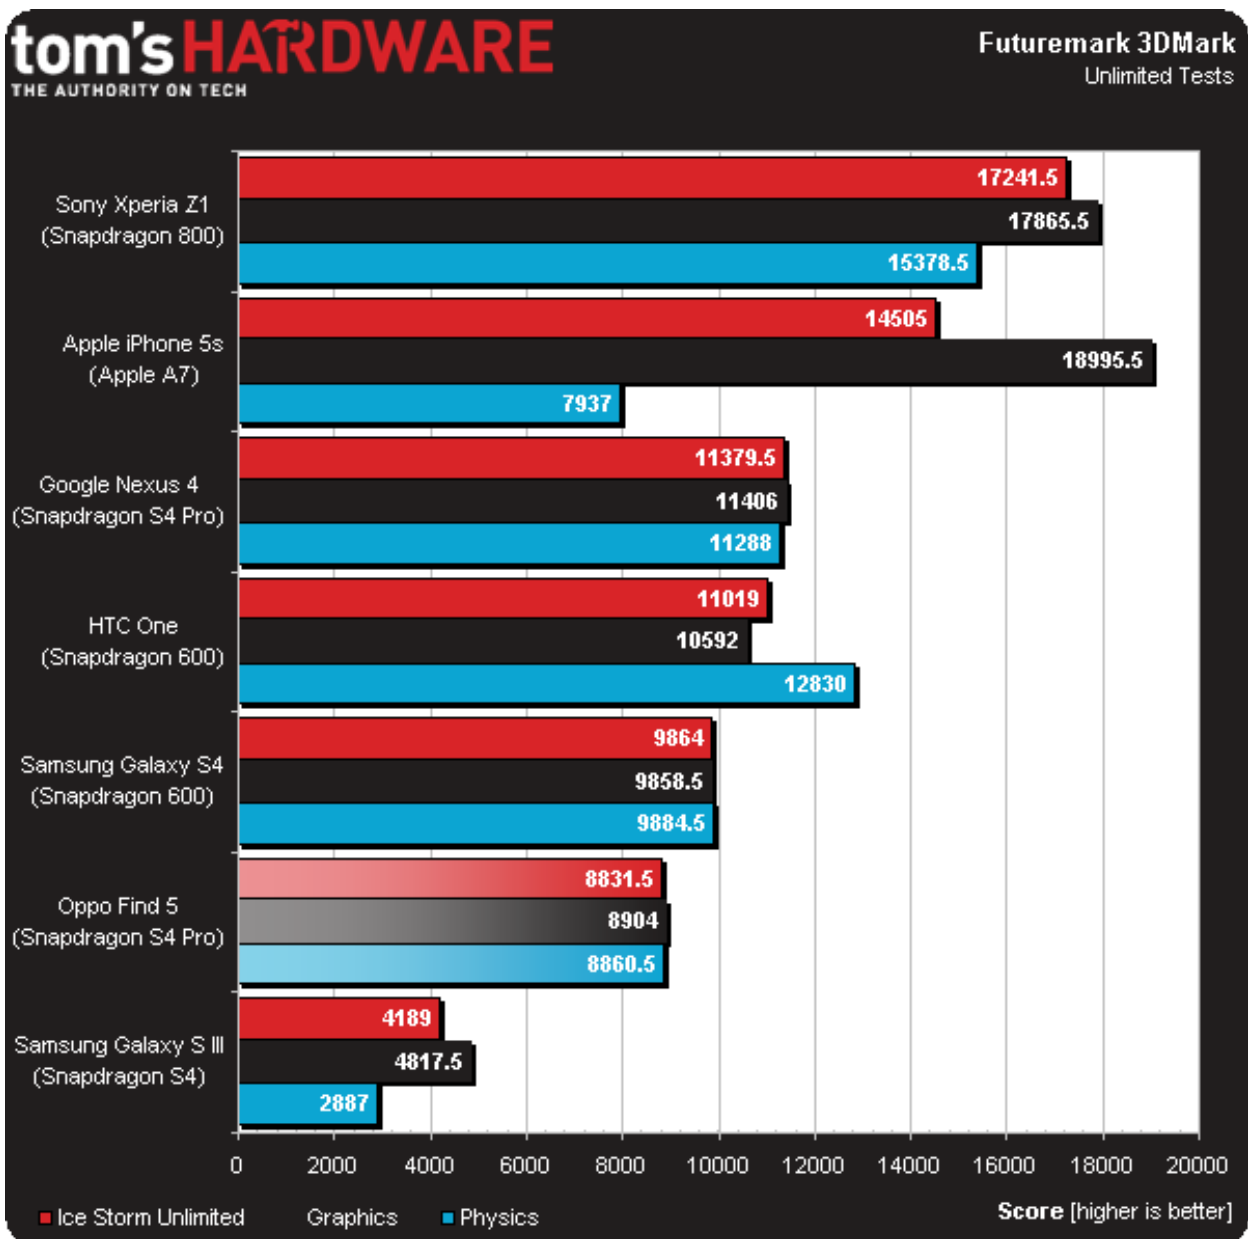


🔥 OpenGL ES 3.0 on iPad Air

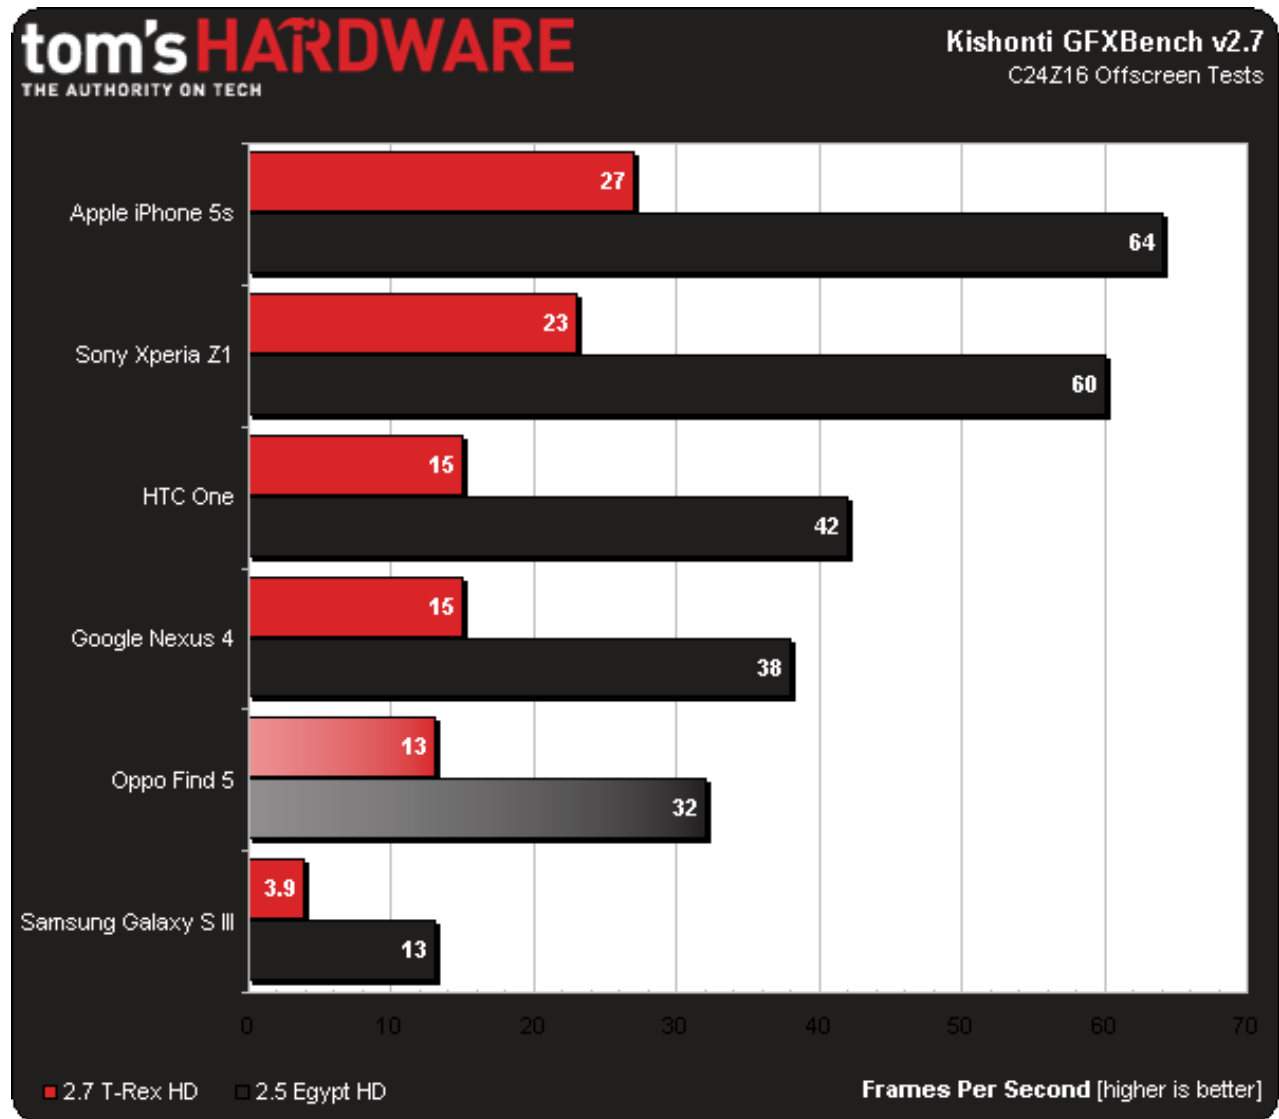


Number of particles in 2D N-body physics simulation

🔥 3DMark results from Tom's Hardware



🔥 GFXBench results from Tom's Hardware



Conclusions

- Still very early days for GPU computing on embedded GPUs
- But it's finally starting to arrive!
- OpenGL ES already quite common
- OpenCL (ES) starting to appear
- Renderscript seems to be lagging
- Availability of dev. kits proving a problem

2014 looks to be the breakout year for fast, easy to use, fully (OpenCL) programmable embedded GPUs!

The screenshot shows a web browser window displaying the website for the microelectronics research group at the University of Bristol. The browser's address bar shows the URL <http://www.cs.bris.ac.uk/Research/Micro/>. The website header includes the University of Bristol logo and the text "microelectronics research group". A navigation menu contains links for HOME, NEWS, DIARY, PUBLICATIONS, COLLABORATION, PROJECTS, PEOPLE, EACO, and WIKI. Below the navigation menu is a large image of a microchip. To the right of the image is a section titled "Upcoming events" with a sub-heading "The Multicore Challenge II: Programming Multicore Systems" and details about a 5 Sep 2011 event at the University of the West of England, Frenchay Campus, Bristol. Below this is a "Supporters & affiliations" section listing logos for Cadence, Mentor Graphics, Infineon, Imagination, Xmos, TVS, ARM, NVIDIA, AMD, and nag. The main content area features three "Recent news" items, each with a sub-heading, date, and a brief description followed by a "read more" link. The first news item is about a research assistant vacancy for massively parallel software libraries for high performance computing, dated 26 Aug 2011. The second is about a research assistant vacancy for adaptive, reliable heterogeneous MPSoCs, dated 24 Aug 2011. The third is about an OpenCL workshop at SC11 to be co-run by Simon McIntosh-Smith, dated 22 Aug 2011. Below the news items is a "Recent publications" section with a sub-heading "Towards Safe Human-Robot Interaction" and a list of authors and the year 2011.


µ Research Group – University of Bristol

Dept. of Computer Science | Dept. of Electronics & Electrical Engineering | Faculty of Engineering | University of Bristol

University of BRISTOL

microelectronics research group

HOME NEWS DIARY PUBLICATIONS COLLABORATION PROJECTS PEOPLE EACO WIKI



Upcoming events

The Multicore Challenge II: Programming Multicore Systems

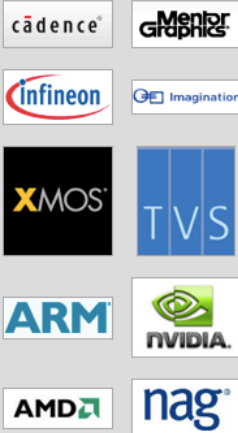
5 Sep 2011 at 1:00 in University of the West of England, Frenchay Campus, Bristol

Experts in multicore technology are coming together in Bristol in September to look at the challenges of developing multicore systems... [read more](#).

[More events...](#)

Supporters & affiliations

The µ Research Group works closely with the following companies and organisations.



Recent news

Research assistant vacancy: massively parallel software libraries for high performance computing

26 Aug 2011
We are looking for another research assistant to work within the group... [read more](#).

Research assistant vacancy: Adaptive, reliable heterogeneous MPSoCs

24 Aug 2011
We are looking for a research assistant to work within the group... [read more](#).

OpenCL workshop at SC11 to be co-run by Simon McIntosh-Smith

22 Aug 2011
Simon McIntosh-Smith will be co-running an all-day workshop at the IEEE/ACM Conference on High Performance Computing, Networking, Storage and Analysis (SuperComputing) with Tim Mattson from Intel and Ben Gaster from AMD... [read more](#).

[Older news...](#)

Recent publications

Towards Safe Human-Robot Interaction

Elena Corina Grigore, [Kerstin Eder](#), Alexander Lenz, Sergey Skachek, Anthony G. Pipe and [Christopher Melhuish](#), 2011