



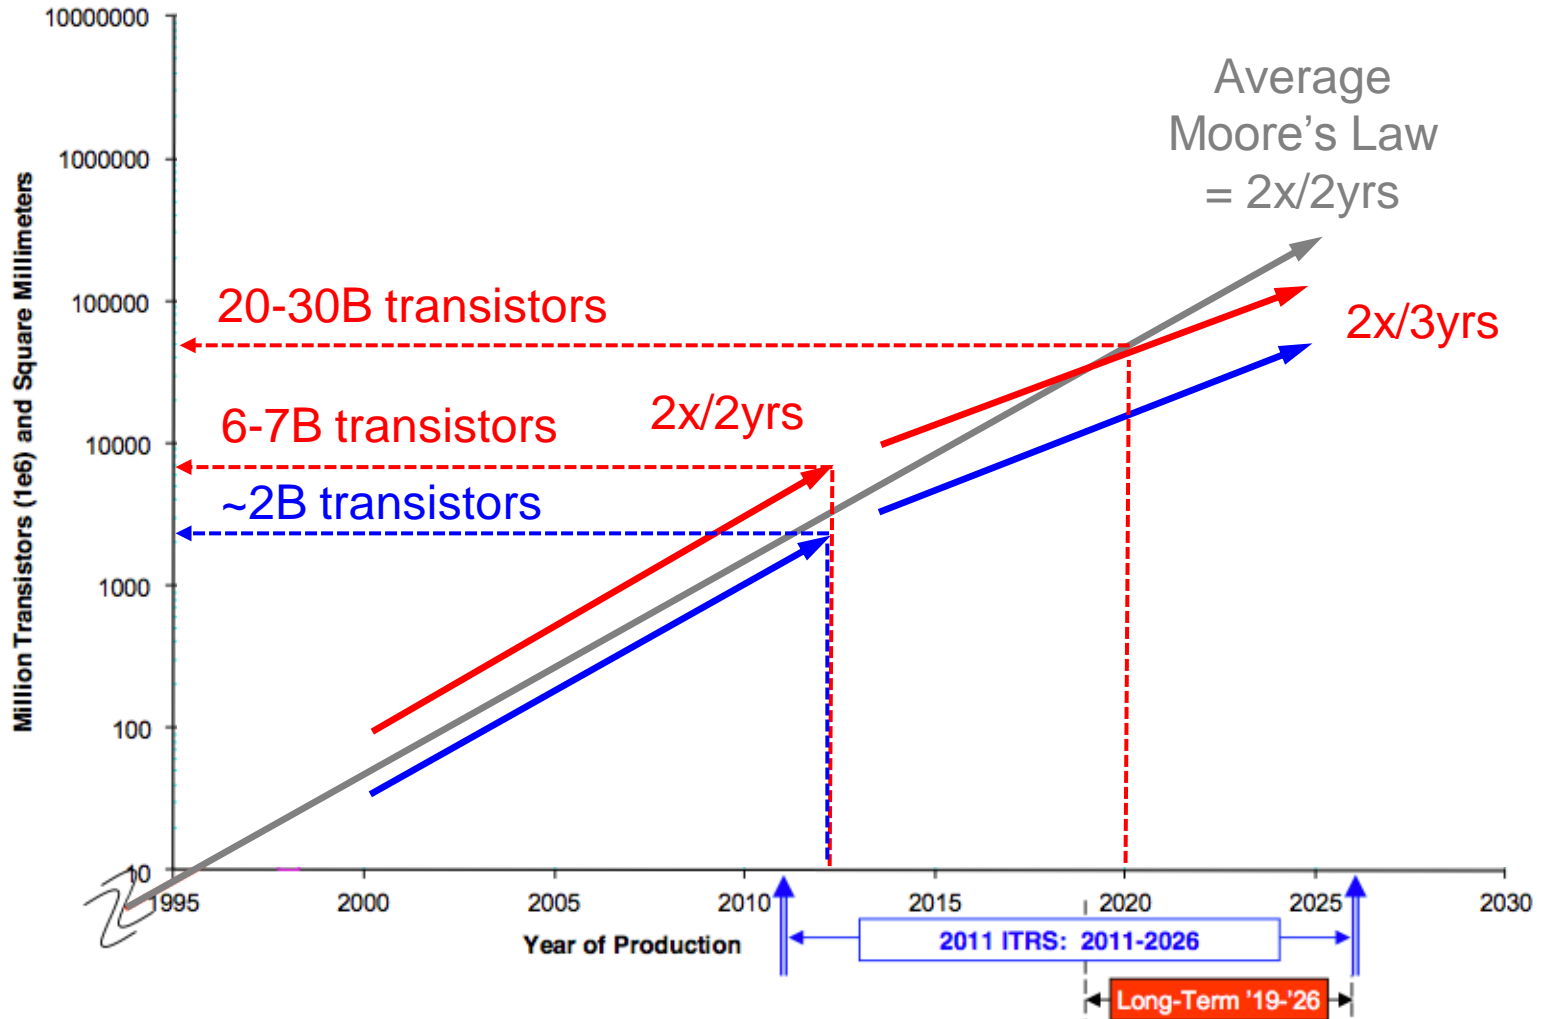
Trends in Heterogeneous Systems Architectures

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




Moore's Law today

2011 ITRS - Functions/chip and Chip Size



🔥 Important technology trends



The real Moore's Law

The clock speed plateau

The power ceiling

Instruction level
parallelism limit

Herb Sutter's latest view

<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.

Causes of heterogeneity

- Multiple types of core
 - CPUs (heavyweight, lightweight, ...)
 - GPUs (massively data parallel)
 - Accelerators (video, encryption, network, ...)
- Interconnect
- Memory type, capacity, ...
- Software (OS, middleware, tools, ...)

Heterogeneous Systems




Intel phi (MIC)



FP7 Mont Blanc ARM + GPU

AMD Fusion APUs



Tegra 3 The World's First Mobile Quad Core, with 5th Companion Core for Low Power

| | |
|---------|--|
| CPU | Quad Core, with 5 th Companion Core — Up to 1.4GHz Single Core, 1.3GHz Quad Core |
| GPU | Up to 3x Higher GPU Performance — 12 Core GeForce GPU |
| VIDEO | Blu-Ray Quality Video — 1080p High Profile @ 40Mbps |
| POWER | Lower Power than Tegra 2 — Variable Symmetric MultiProcessing (vSMP) |
| MEMORY | Up to 3x Higher Memory Bandwidth — DDR3L-1500, LPDDR2-1066 |
| IMAGING | Up to 2x Faster ISP (Image Signal Processor) |
| AUDIO | HD Audio, 7.1 channel surround |
| STORAGE | 2-6x Faster — MMC 4.41, SD3.0, SATA-II |

NVIDIA Tegra, Project Denver

🔥 Heterogeneity is mainstream



Quad-core ARM Cortex A9 CPU

Quad-core SGX543MP4+ Imagination GPU

🔥 Implications for software

- New programming languages, models, ...
 - OpenCL is designed specifically for this
- Dynamically adaptive software
 - Discover resources at run-time
 - Load balancing
 - Auto-tuning
- Application frameworks and libraries



Conclusions

- Heterogeneity is an increasingly important trend
 - An inevitable companion to increasing parallelism?
- The market is finally starting to create and adopt the necessary open standards ([OpenCL](#), [HSA](#), ...)
- Parallel programming models are likely to (re)proliferate
- Exciting times ahead!