

GPUs and Accelerators in HPC



Knights Ferry

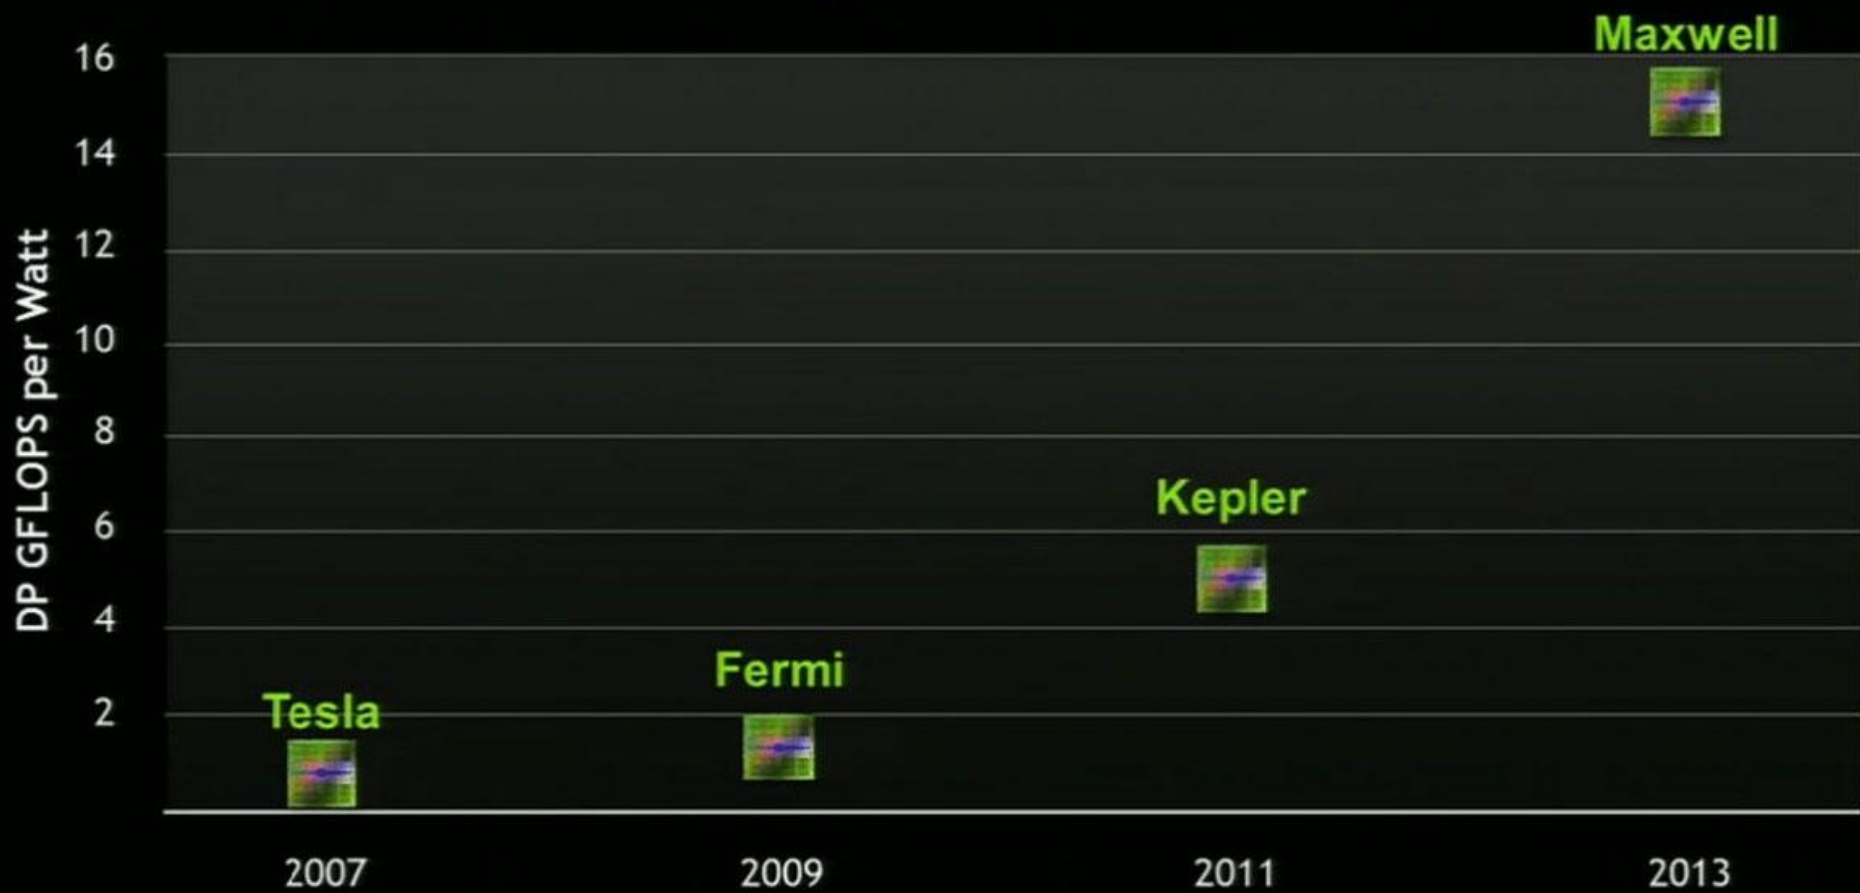


- Software development platform
- Growing availability through 2010
- 32 cores, 1.2 GHz
- 128 threads at 4 threads / core
- 8MB shared coherent cache
- 1-2GB GDDR5
- Bundled with Intel HPC tools

Software development platform for Intel® MIC architecture



CUDA GPU Roadmap




Daresbury resources for accelerated computing

- IBM PowerXCell 8i cluster
- NVIDIA Tesla S1070
- AMD/ATI FireStream 9270
- NVIDIA Fermi – expecting soon

cseht

- SuperMicro X8DAi m/b
- 2x Intel Xeon E5540
- 24GB DDR3 per node
- 8x Tesla S1070
- Voltaire HCA410-4EX

~ 70 users



GPU Programming

- Libraries and frameworks
 - cuBLAS, cuFFT
 - CULA
- Directives
 - HMPP
 - PGI accelerator model
 - PathScale
- API
 - C for CUDA, CUDA Fortran
 - OpenCL
 - DirectCompute



Software

- CUDA 3.1 (updating to 3.2 shortly)
- ATI Stream SDK 2.2
- HMPP 2.3.1
- PGI 10.8 with accelerator support
- PathScale ENZO compiler
- cuBLAS, cuFFT, CULA 2.0, ACML
- Intel compilers and tools, OpenMPI 1.3 etc

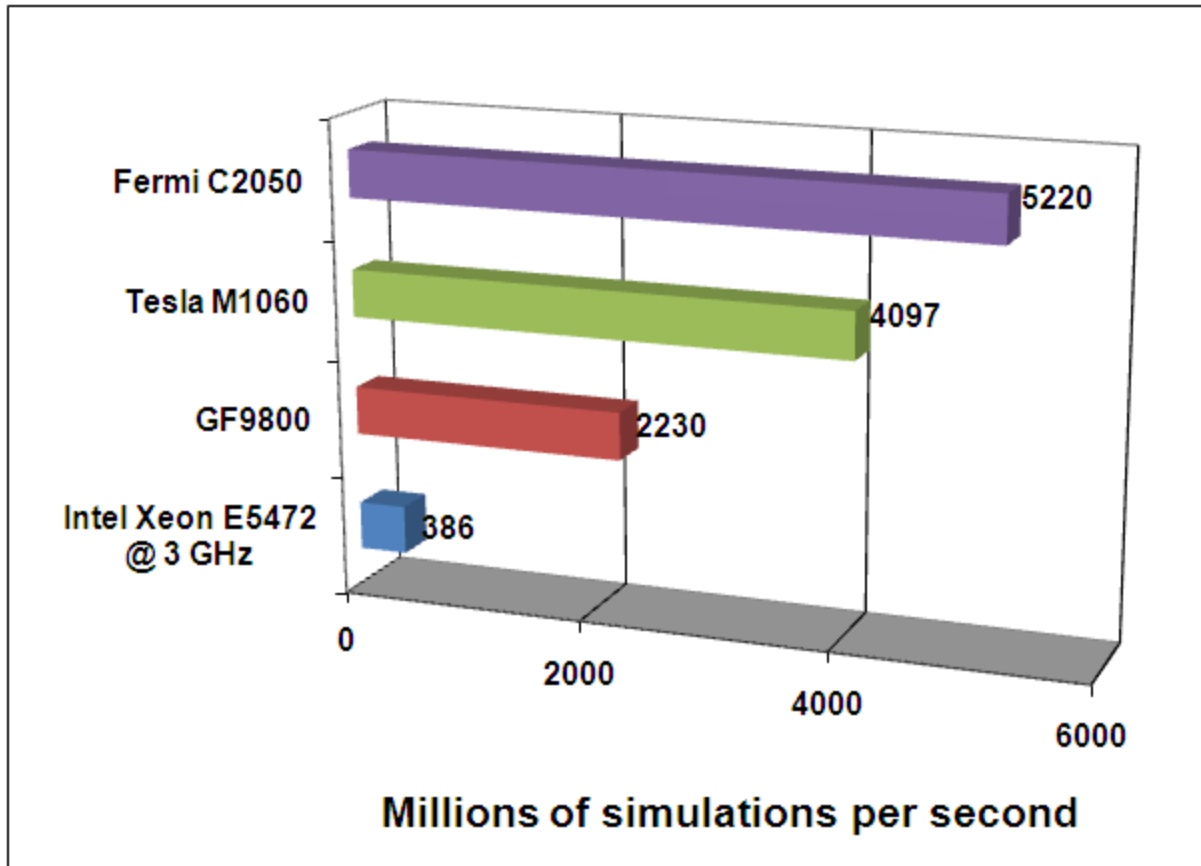
More details <http://www.cse.scitech.ac.uk/disco/cseht/cseht.shtml>

Apply for a free account <http://www.cse.scitech.ac.uk/disco/useraccounts.shtml>

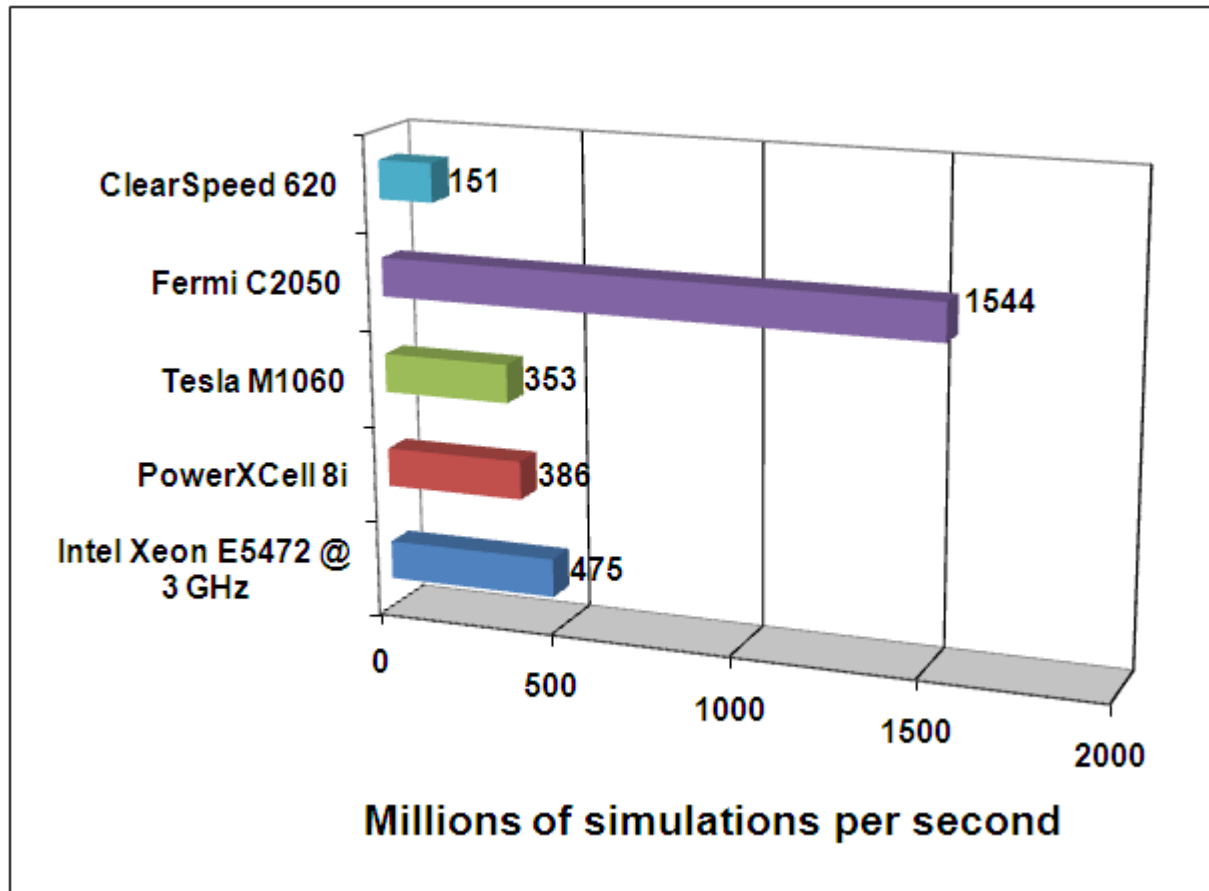
GPU related projects

- **G2X – Gadget2 port to CUDA**
Carsten Frigaard - STFC, Durham, NVIDIA
- **DLPOLY port to CUDA**
Christos Kartsaklis, Ruari Nestor - ICHEC, STFC
- **GAMESS-UK**
Karl Wilkinson - Cape Town & STFC
- **GNEMO – GPU acceleration of the NEMO Ocean Model**
Andy Porter - STFC
- **Power efficiency**
Igor Kozin - STFC

European option pricing - SP



European option pricing - DP



CUDA vs HMPP

(European option pricing)

	single precision			double precision		
	CUDA	HMPP		CUDA	HMPP	
		CUDA	OpenCL		CUDA	OpenCL
M1060	4100	3080	2490	350	370	250
C2050	4720	3290	2510	1450	1440	1010
FS9270			126			n/a