




TESLA

GPU Computing



NVIDIA Tesla C2050 Performance Benchmarks
April 2010

Tesla C-Series Workstation GPUs




NVIDIA

	Tesla C1060	Tesla C2050	Tesla C2070
Architecture	Tesla 10-series GPU	Tesla 20-series GPU	
Number of Cores	240	448	
Caches	16 KB Shared Memory / 8 cores	64 KB L1 cache + Shared Memory / 32 cores, 768 KB L2 cache	
FP Peak Performance	933 GFlops (single) 78 GFlops (double)	1,03 TFlops (single) 515 GFlops (double)	
FP Application Efficiency (Tesla C1060 reference)	1 (single) 1 (double)	1.5 - 2 (single) 3 - 4 (double)	
GPU Memory	4 GB GDDR3	3 GB GDDR5 2.625 GB with ECC on	6 GB GDDR5 5.25 GB with ECC on
Memory Bandwidth	102 GB/s	144 GB/s ECC off 115 GB/s ECC on	
Video Output	No	DVI-I	
System I/O	PCIe x16 Gen2 (single way async. transfer)	PCIe x16 Gen2 (dual ways async. transfer)	
Power	188 W max	247 W max	225 W max



May 2010

Summer 2010

Performance Summary



Standard FFT Library: cuFFT 3.1-Pre-release




cuFFT 3.1-pre-release: NVIDIA Tesla C1060 GPU and Tesla C2050 (Fermi)
 MKL 10.1r1: Quad-Core Intel Core i7 (Nehalem) 3.2GHz


Standard BLAS Library: cuBLAS 3.1-Pre-release



GFlops Single Precision BLAS: SGEMM



GFlops Double Precision BLAS: DGEMM



cuBLAS: CUDA 3.1-pre-release: , Tesla C1060 and Tesla C2050

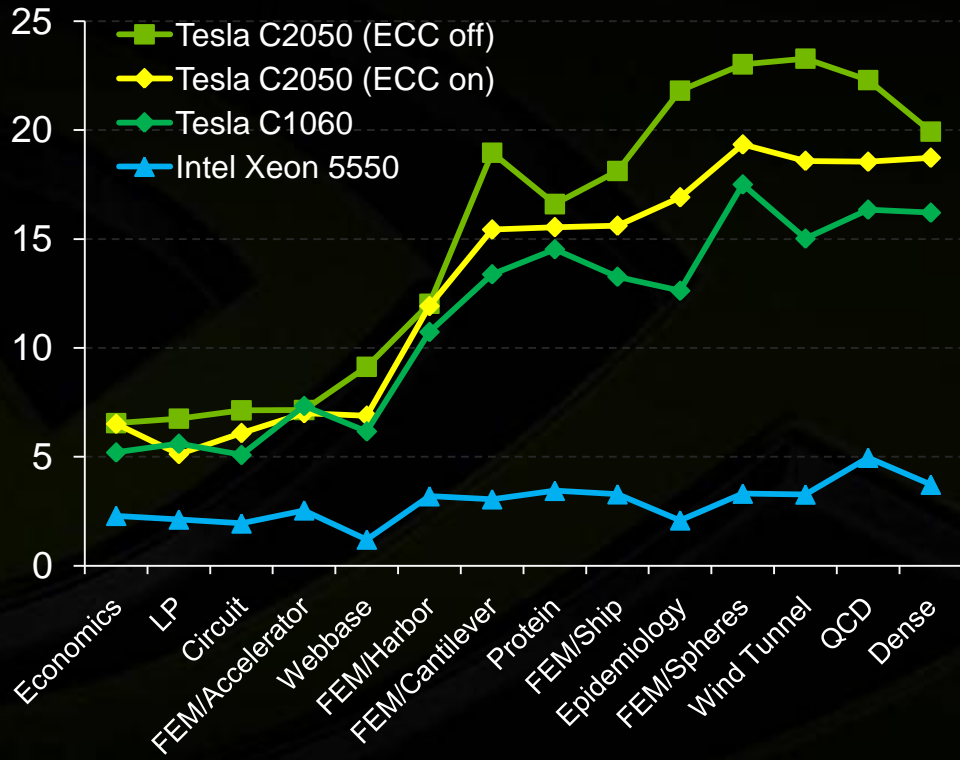
MKL 10.1r1: Intel Core2 Extreme, 3.00GHz

Sparse Matrix-Vector Multiplication (SpMV)



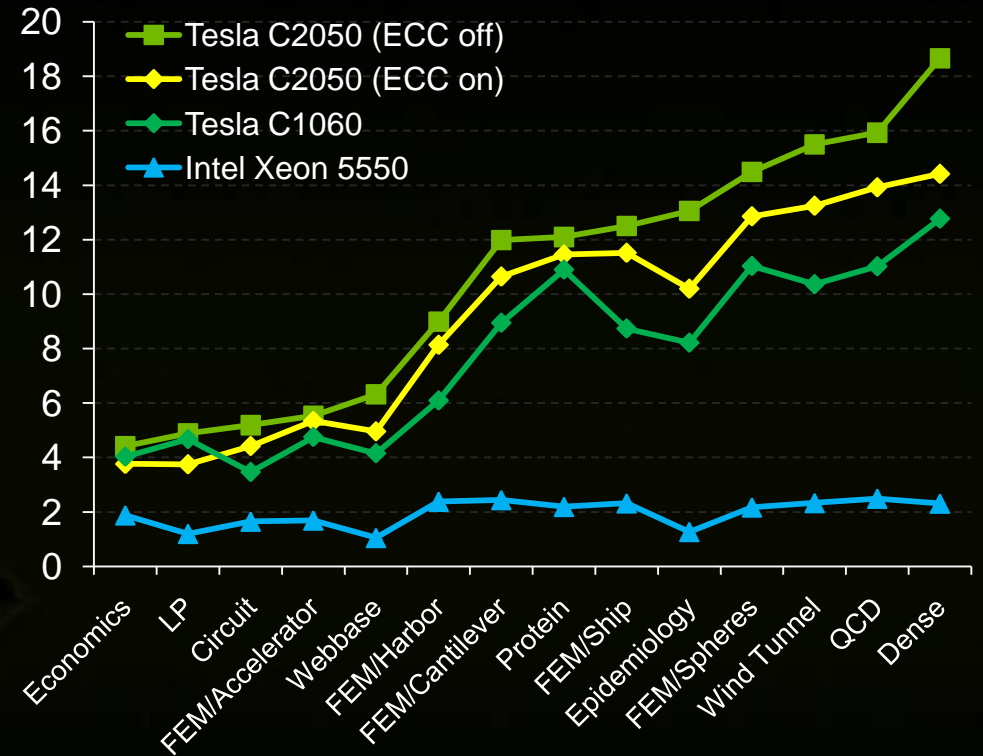
GFlops

Single Precision



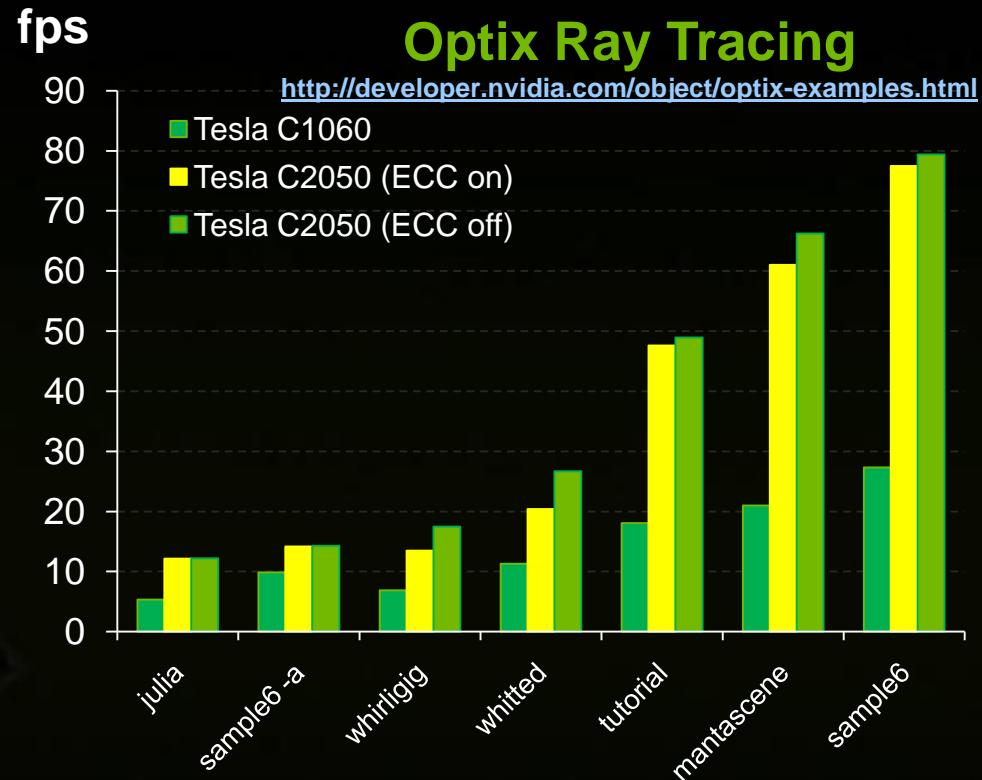

GFlops

Double Precision



SpMv: CUDA 3.0, Tesla C1060 and Tesla C2050
MKL 10.2: Intel Xeon 5550, 2.67 GHz

N-Body and Ray Tracing Performance




CUDA 3.0, Tesla C1060 and Tesla C2050

8x Higher Linpack


8x

Performance
GFlops




6x

Performance / \$
GFlops / \$K



5x

Performance / Watt
GFlops / KWatt



CPU 1U Server: 2x Intel Xeon X5550 (Nehalem) 2.66 GHz, 48 GB memory, \$7K, 0.67 KW

GPU-CPU 1U Server: 2x Tesla C2050 + 2x Intel Xeon X5550, 48 GB memory, \$11K, 1.0 KW

“In testing our key applications, the Tesla GPUs delivered speed-ups that we had never seen before, sometimes even orders of magnitude.”



Satoshi Matsuoka

Professor
Tokyo Institute of Technology

“Future computing architectures will be hybrid systems with parallel-core GPUs working in tandem with multi-core CPUs”

Jack Dongarra
Professor, University of Tennessee
Author of Linpack

