

# New NVIDIA Tesla GPUs Reduce Cost Of Supercomputing By A Factor Of 10

November 16 2009

---



Tesla C2050 GPU Computing Processor

NVIDIA Corporation today unveiled the Tesla 20-series of parallel processors for the high performance computing (HPC) market, based on its new generation CUDA processor architecture, codenamed "Fermi".

Designed from the ground-up for [parallel computing](#), the [NVIDIA](#) Tesla 20-series GPUs slash the cost of computing by delivering the same performance of a traditional CPU-based cluster at one-tenth the cost and one-twentieth the power.

The Tesla 20-series introduces features that enable many new applications to perform dramatically faster using [GPU](#) Computing.

These include ray tracing, 3D cloud computing, video encoding, database search, data analytics, computer-aided engineering and virus scanning.

“NVIDIA has deployed a highly attractive architecture in Fermi, with a feature set that opens the technology up to the entire computing industry,” said Jack Dongarra, director of the Innovative Computing Laboratory at the University of Tennessee and co-author of LINPACK and LAPACK.



The Tesla 20-series GPUs combine parallel computing features that have never been offered on a single device before. These include:

- Support for the next generation IEEE 754-2008 double precision floating point standard
- ECC (error correcting codes) for uncompromised reliability and accuracy
- Multi-level cache hierarchy with L1 and L2 caches
- Support for the C++ programming language
- Up to 1 terabyte of memory, concurrent kernel execution, fast context switching, 10x faster atomic instructions, 64-bit virtual address space, system calls and recursive functions

At their core, Tesla GPUs are based on the massively parallel CUDA computing architecture that offers developers a parallel computing model that is easier to understand and program than any of the alternatives developed over the last 50 years.

"There can be no doubt that the future of computing is parallel processing, and it is vital that computer science students get a solid grounding in how to program new parallel architectures," said Dr. Wen-mei Hwu, Professor in Electrical and Computer Engineering of the University of Illinois at Urbana-Champaign. "GPUs and the CUDA programming model enable students to quickly understand parallel programming concepts and immediately get transformative speed increases."

The family of Tesla 20-series GPUs includes:

1. Tesla C2050 & C2070 GPU Computing Processors

- Single GPU PCI-Express Gen-2 cards for workstation configurations
- Up to 3GB and 6GB (respectively) on-board GDDR5 memoryi
- Double precision performance in the range of 520GFlops - 630 GFlops

2. Tesla S2050 & S2070 GPU Computing Systems

- Four Tesla GPUs in a 1U system product for cluster and datacenter deployments
- Up to 12 GB and 24 GB (respectively) total system memory on board GDDR5 memoryii
- Double precision performance in the range of 2.1 TFlops - 2.5 TFlops

The Tesla C2050 and C2070 products will retail for \$2,499 and \$3,999 and the Tesla S2050 and S2070 will retail for \$12,995 and \$18,995. Products will be available in Q2 2010.

Source: NVIDIA

Citation: New NVIDIA Tesla GPUs Reduce Cost Of Supercomputing By A Factor Of 10 (2009, November 16) retrieved 26 April 2024 from <https://phys.org/news/2009-11-nvidia-tesla-gpus-supercomputing-factor.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.