

TESLA™ C2050 / C2070 GPU COMPUTING PROCESSOR SUPERCOMPUTING AT 1/10TH THE COST

The NVIDIA Tesla™ C2050 and C2070 Computing Processors fuel the transition to parallel computing and bring the performance of a small cluster to the desktop.

Based on the next-generation CUDA™ architecture codenamed “Fermi”, the 20-series family of Tesla GPUs support many “must have” features for technical and enterprise computing including C++ support, ECC memory for uncompromised accuracy and scalability, and a 7X increase in double precision performance compared Tesla 10-series GPUs. The Tesla C2050 and C2070 GPUs are designed to redefine high performance computing and make supercomputing available to everyone.


Compared to the latest quad-core CPUs, Tesla C2050 and C2070 Computing Processors deliver equivalent supercomputing performance at 1/10th the cost and 1/20th the power consumption.



BIOSCIENCE



FINANCE



MEDICAL

TECHNICAL SPECIFICATIONS

- FORM FACTOR
 - > 9.75" PCIe x16 form factor
- # OF TESLA GPUS
 - > 1
- # OF CUDA CORES
 - > 448
- FREQUENCY OF CUDA CORES
 - > 1.15 GHz
- DOUBLE PRECISION FLOATING POINT PERFORMANCE (PEAK)
 - > 515 Gflops
- SINGLE PRECISION FLOATING POINT PERFORMANCE (PEAK)
 - > 1.03 Tflops
- TOTAL DEDICATED MEMORY
 - > Tesla C2050: 3GB GDDR5*
 - > Tesla C2070: 6GB GDDR5*
- MEMORY SPEED
 - > 1.55 GHz
- MEMORY INTERFACE
 - > 384-bit
- MEMORY BANDWIDTH
 - > 148 GB/sec
- POWER CONSUMPTION
 - > 225W TDP
- SYSTEM INTERFACE
 - > PCIe x16 Gen2
- THERMAL SOLUTION
 - > Active Fansink
- SOFTWARE DEVELOPMENT TOOLS
 - > CUDA C/C++/Fortran, OpenCL, DirectCompute Toolkits, NVIDIA Parallel Nsight™ for Visual Studio

*Note: With ECC on, a portion of the dedicated memory is used for ECC bits, so the available user memory is reduced by 12.5%. [e.g. 3 GB total memory yields 2.625 GB of user available memory.]

FEATURES AND BENEFITS

GPUS POWERED BY THE FERMI-GENERATION OF THE CUDA ARCHITECTURE	Delivers cluster performance at 1/10 th the cost and 1/20 th the power of CPU-only systems based on the latest quad core CPUs.
448 CUDA CORES	Delivers up to 515 Gigaflops of double-precision peak performance in each GPU, enabling a single workstation to deliver a Teraflop or more of performance. Single precision peak performance is over a Teraflop per GPU.
ECC MEMORY	Meets a critical requirement for computing accuracy and reliability for workstations. Offers protection of data in memory to enhance data integrity and reliability for applications. Register files, L1/L2 caches, shared memory, and DRAM all are ECC protected.
DESKTOP CLUSTER PERFORMANCE	Solves large-scale problems faster than a small server cluster on a single workstation with multiple GPUs.
UP TO 6GB OF GDDR5 MEMORY PER GPU	Maximizes performance and reduces data transfers by keeping larger data sets in local memory that is attached directly to the GPU.
NVIDIA PARALLEL DATACACHE™	Accelerates algorithms such as physics solvers, ray-tracing, and sparse matrix multiplication where data addresses are not known beforehand. This includes a configurable L1 cache per Streaming Multiprocessor block and a unified L2 cache for all of the processor cores.
NVIDIA GIGATHREAD™ ENGINE	Maximizes the throughput by faster context switching that is 10X faster than previous architecture, concurrent kernel execution, and improved thread block scheduling.
ASYNCHRONOUS TRANSFER	Turbocharges system performance by transferring data over the PCIe bus while the computing cores are crunching other data. Even applications with heavy data-transfer requirements, such as seismic processing, can maximize the computing efficiency by transferring data to local memory before it is needed.
CUDA PROGRAMMING ENVIRONMENT WITH BROAD SUPPORT OF PROGRAMMING LANGUAGES AND APIs	Choose C, C++, OpenCL, DirectCompute, or Fortran to express application parallelism and take advantage of the “Fermi” GPU’s innovative architecture. NVIDIA Parallel Nsight tool is available for Microsoft Visual Studio developers.
HIGH SPEED, PCIe GEN 2.0 DATA TRANSFER	Maximizes bandwidth between the host system and the Tesla processors. Enables Tesla systems to work with virtually any PCIe-compliant host system with an open PCIe x16 slot.

DRIVERS AND DOWNLOADS

- > Tesla C2050/C2070 is supported under
 - Linux 32-bit and 64-bit
 - Windows XP, Windows Vista, Windows 7 (32-bit and 64-bit)
- > Vertical Solutions and Software page: www.nvidia.com/object/vertical_solutions.html
- > Software for Tesla C2050 / C2070
 - Download latest driver at www.nvidia.com/Download/index.aspx?lang=en-us
 - Software development tools is available at www.nvidia.com/object/tesla_software.html

SUPPORT

- > **HARDWARE SUPPORT**
For RMA requests, replacements and warranty issues regarding your NVIDIA based product, please contact the reseller that you purchased it from.
- > **KNOWLEDGEBASE**
Our knowledgebase is available online 24x7x365 and contains answers to the most common questions and issues: www.nvidia.custhelp.com/cgi-bin/nvidia.cfg/php/enduser/std_alp.php
- > **USER FORUMS**
Discuss Tesla products, talk about CUDA development, and share interesting issues, tips and solutions with your fellow NVIDIA Tesla users on the CUDA discussion forums: www.forums.nvidia.com
- > **PRE-PURCHASING SUPPORT**
Pre-purchasing support is available to assist you in deciding which Tesla unit to purchase for your application: www.nvidia.com/object/tesla_assistance.html

To learn more about NVIDIA Tesla, go to www.nvidia.com/tesla