




CUDA

Ian Buck

The complexity of the problem is hidden
by the simplicity of the solution.

CUDA Development

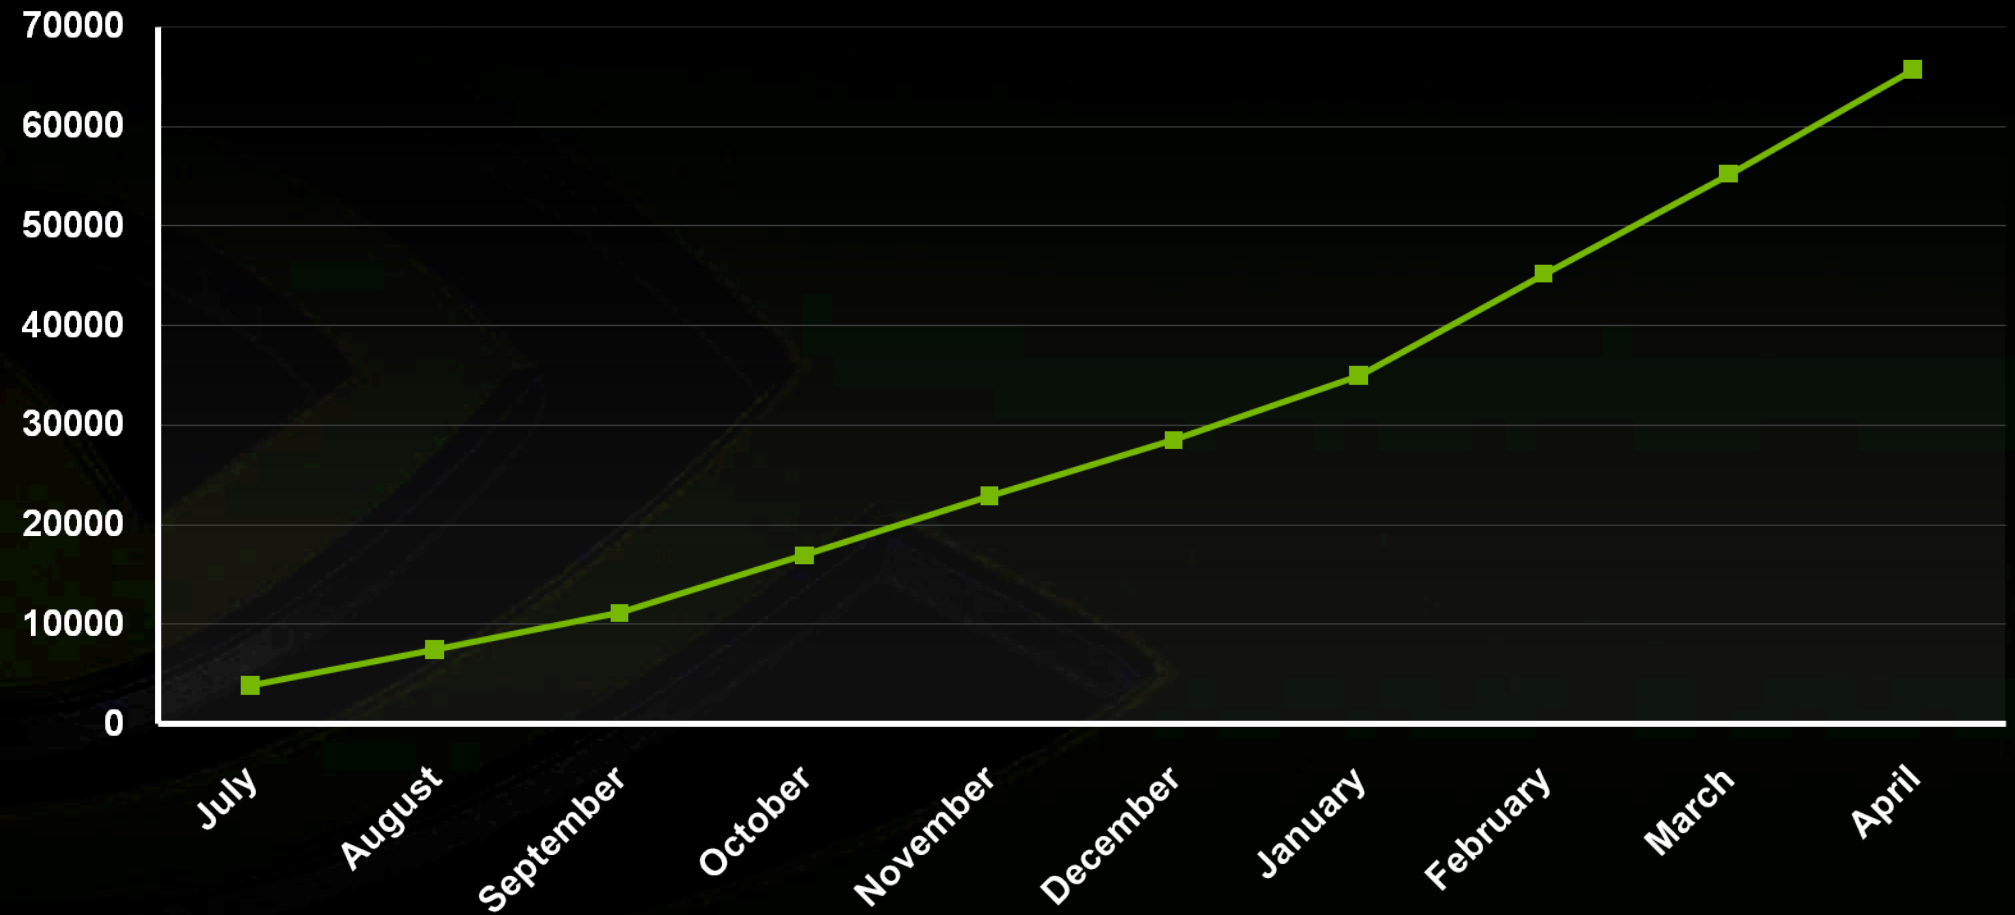


CUDA public beta

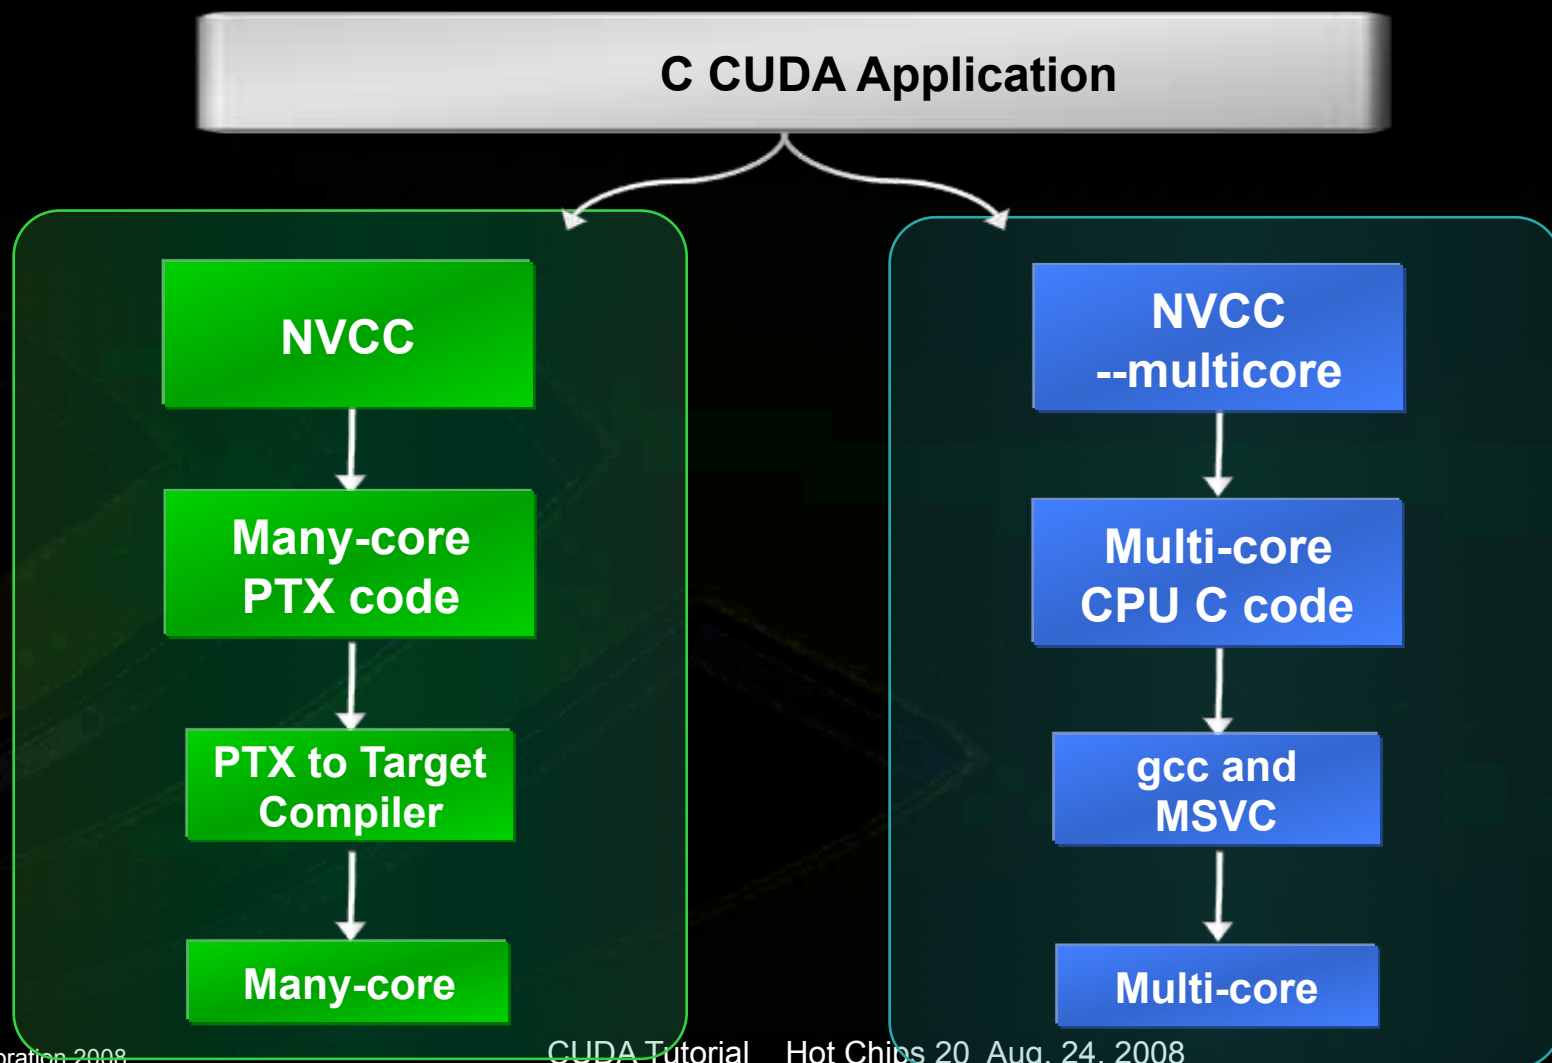
CUDA in production

CUDA 2.0

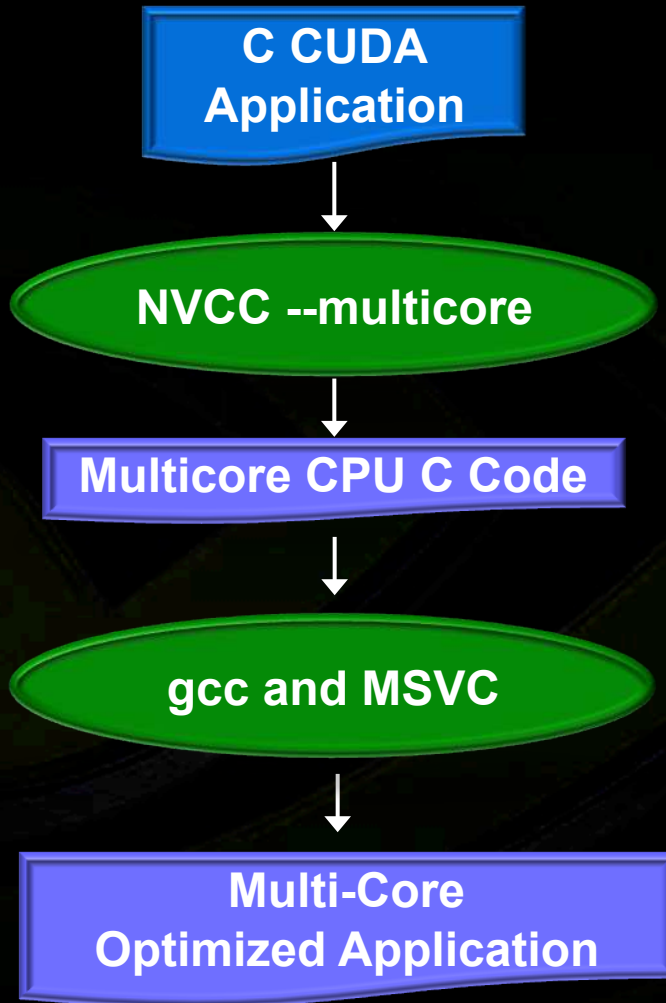
CUDA Compiler Downloads



CUDA 2.0: Many-core + Multi-core support



Compiling CUDA for Multi-Core



- **NVCC generates code for multi-core**
- **Alpha in July**
- **More cores = more performance**

What's Next for CUDA

Fortran

C++

Multiple GPUs

Debugger

Profiler

GPU Cluster

Building a 100TF datacenter

CPU 1U Server



4 CPU cores

0.07 Teraflop

\$ 2000

400 W

1429 CPU servers

\$ 3.1 M

571 KW

4 GPUs: 960 cores

4 Teraflops

\$ 8000

700 W

25 CPU servers
25 Tesla systems

\$ 0.31 M


27 KW

Tesla 1U System



10x lower cost

21x lower power

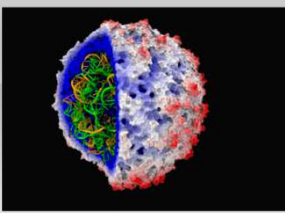


Beyond Gaming



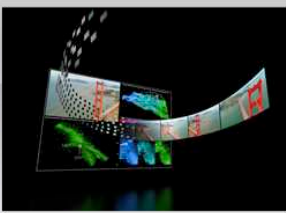
146X

Interactive visualization of volumetric white matter connectivity



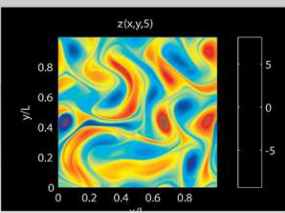
36X

Ionic placement for molecular dynamics simulation on GPU



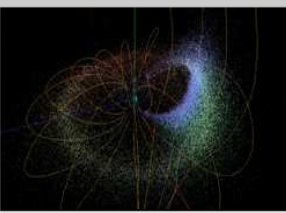
11X

Transcoding HD video stream to H.264 for portable video



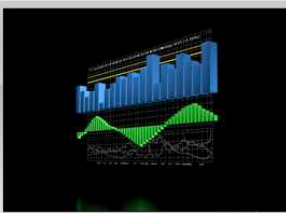
17X

Simulation in Matlab using .mex file CUDA function



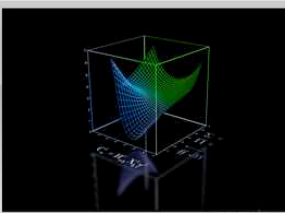
100X

Astrophysics N-body simulation




149X

Financial simulation of LIBOR model with swaptions



47X

GLAME@lab: An M-script API for linear Algebra operations on GPU



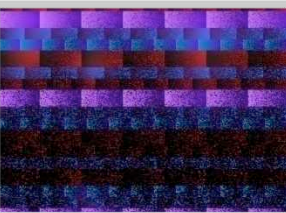
20X

Ultrasound medical imaging for cancer diagnostics



24X

Highly optimized object oriented molecular dynamics



30X

Cmatch exact string matching to find similar proteins and gene sequences

Design: CAD Design For Apparel Cloth Physics




CUDA Accelerates MATLAB[®]

17x Faster with GPU Compute




Pseudo-spectral simulation of 2D Isotropic turbulence



1024x1024 mesh, 400 RK4 steps, Windows XP, Core2 Duo 2.4Ghz vs GeForce 8800GTX


http://www.amath.washington.edu/courses/571-winter-2006/matlab/FS_2Dturb.m



CUDA Advantage on MATLAB® Image Toolbox



MATLAB Image Deblur using CUDA
1024 x 1024




Dual Xeon Core 2 Quad 5320 1.86GHz (8 cores) 4GB
NVIDIA Quadro 5600 1.5GB


Using MATLAB Lucy-Richardson Deconvolution: `deconvlucy(. . .)`
Results have max pixel difference of no more than 1 grey level




LATEST CUDA NEWS Calling All C Programming Gurus – Signup for the CUDA Coding Contest and Enter to Win Some Cash




Efficient Computation of Sum Products on GPUs 270 x



Molecular Dynamics Simulations on GPUs 150 x




Fast k Nearest Neighbor Search using GPU 120 x




Prestack Seismic Data Interaction 100 x




Visual Molecular Dynamics: VMD 100 x




Accelerating Statistical Static Timing Analysis 260 x



Fast Support Vector Machine Training and Classification 138 x




Interactive Visualization of Volumetric White Matter Connectivity 100 x



Teraflop CFD Computing 100 x




SciFinance Speeds Financial Results with Parallel Computing 80 x




Map Reduce Framework 150 x




Quantum Chemistry Two-Electron Integral Evolution 130 x



High Performance Pattern Recognition on GPU 100 x



Geometric Algorithms with CUDA 100 x



Simulation Open Framework Architecture (SOFA) 55 x

Search Sort by Speed Up Share Your Work




Search Sort by Speed Up Share Your Work

Search Sort by Speed Up Share Your Work

Map Reduce Framework 120 x Efficient Computation of Sum Products on GPU 130 x Recognition on GPU 100 x CFD Architecture (SOFA) 55 x




70M CUDA GPUs




Heterogeneous Computing


CUDA




Oil & Gas




Finance



Medical



Biophysics



Numerics



Audio



Video



Imaging