TESLA

GPU Computing

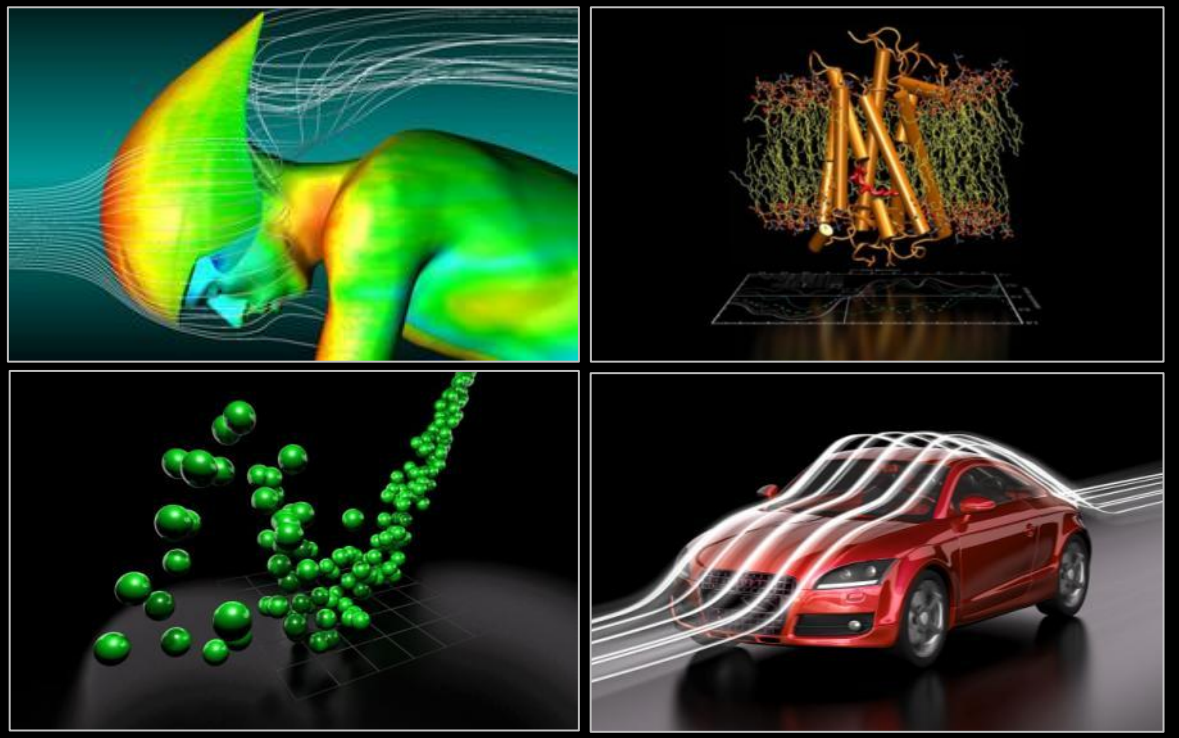

GPU Accelerated Applications for Academic Research

Devang Sachdev
Sr. Product Manager

Explosive Growth of GPU Computing

8,000 Institutions with CUDA Developers

1,500,000 CUDA Downloads

395,000,000 CUDA GPUs Shipped

629 University Courses
In 62 Countries

Titan: World's #1 Open Science Supercomputer

18,688 Tesla Kepler GPUs

Fully GPU Accelerated Computing System

20+ Petaflops: 90% of Performance from GPUs

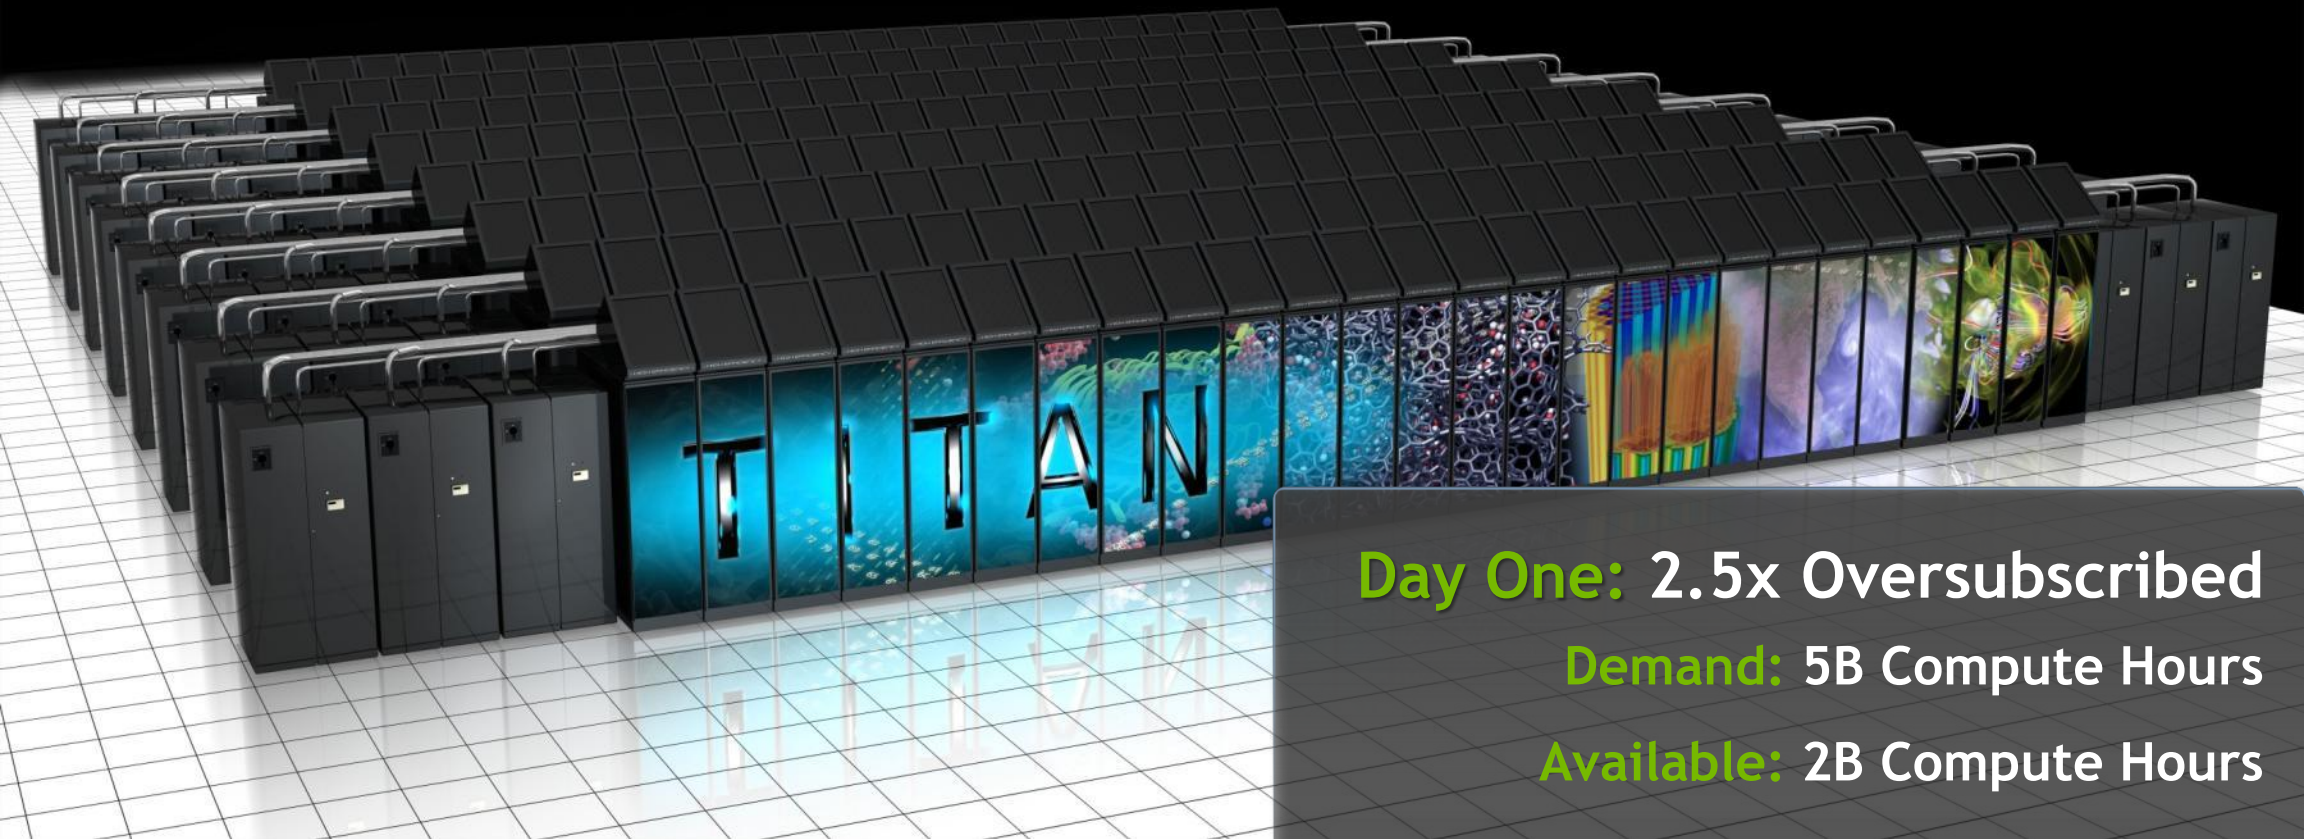

Day One: 2.5x Oversubscribed

Demand: 5B Compute Hours

Available: 2B Compute Hours

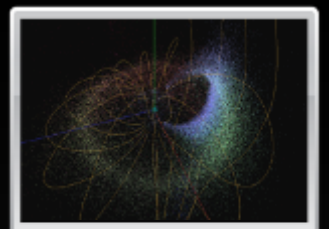

100X

Astrophysics
RIKEN

130X

Quantum Chemistry
U of Illinois, Urbana

25X

Fast Flow Simulation
FluiDyna

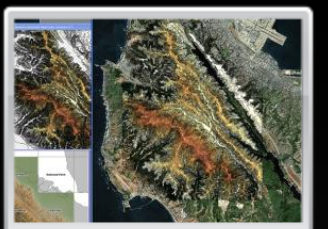

60X

Geospatial Imagery
PCI Geomatics

GPUs Accelerate HPC

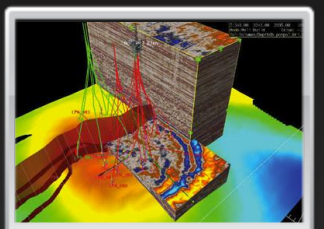

12X

Kirchoff Time Migration
Acceleware

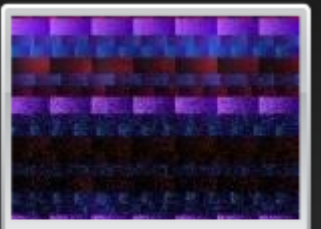

30X

Gene Sequencing
U of Maryland

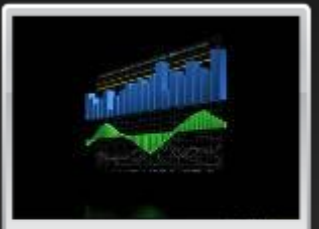

149X

Financial Simulation
Oxford

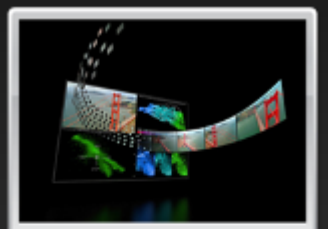

18X

Video Transcoding
Elemental Tech

GPUs Accelerate Compute-Intensive Functions

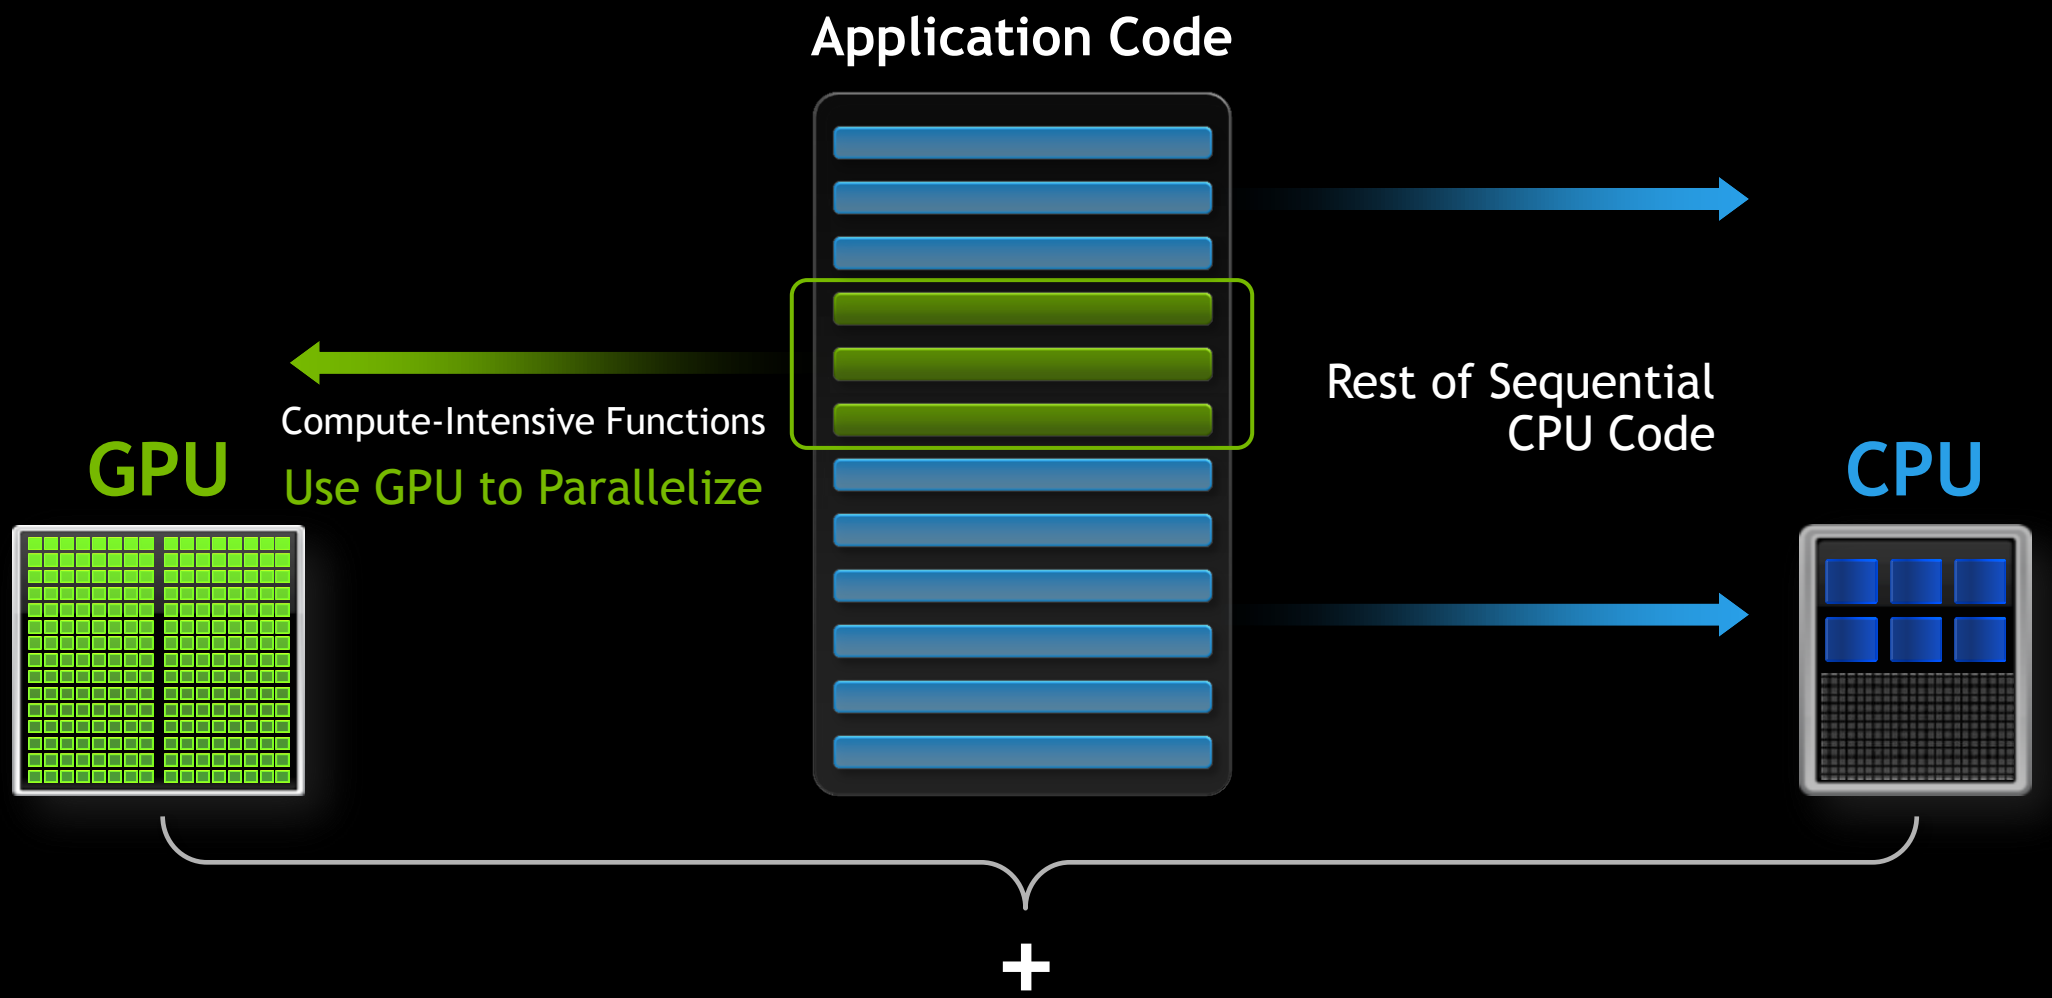

GPU Computing - Easier Than You Think

GPU Apps	Libraries	Directives	CUDA
NAMD	C++ Thrust	OpenACC	CUDA C
MATLAB	cuBLAS	*Open*	CUDA C++
CHROMA	cuSPARSE	*Simple*	CUDA Fortran
VASP	cuFFT	*Portable*	
Many more	*Many more*		

Key Application Developers Embracing GPUs

"GPU computing enables highly accurate calculations with a short time to solution."

Dr. Jeongnim Kim, Lead Developer for QMCPACK

"GPU capabilities are the path forward for scaling dense codes, and VASP is no exception."

Dr. Michael Widom,
Professor at Carnegie Mellon University &
One of the leads in the VASP++ Consortium

"A petascale GPU cluster would dramatically increase access to large-scale simulation capability in the 10M-100M-atom range, allowing simulations of significant cellular and viral structures on useful timescales."

Dr. James Philips, Lead Developer for NAMD

GPU Apps Grow Over 60% in 2012

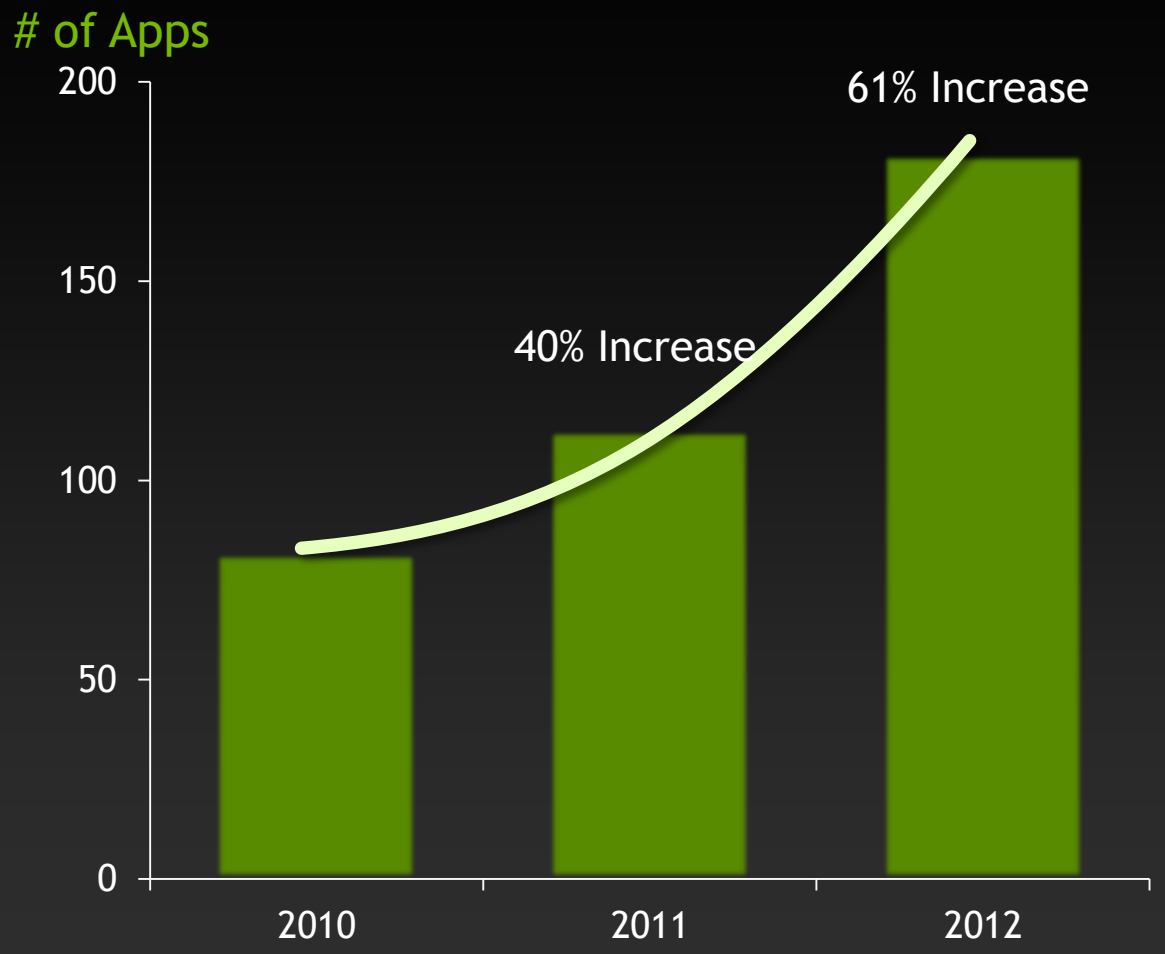

- Astrophysics
- Computational Biology
- Computational Chemistry
- Defense & Intelligence
- Electronics Design Automation
- Energy
- Finance
- Fluid Dynamics
- High Energy & Nuclear Physics
- Industrial Design
- Media & Entertainment
- Nuclear & Fusion Energy
- Structural Mechanics
- Weather/Climate

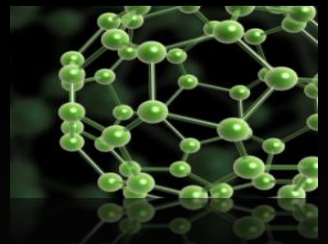

Computational Chemistry

AMBER

NAMD

LAMMPS

GROMACS

Quantum Espresso

DL_PLOY

GAMESS

NWChem

Q-CHEM

TeraChem

Abalone

ACEMD

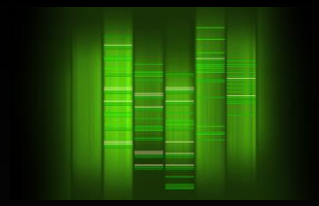

Bioinfomatics

BarraCUDA

CUDASW++

GPU-BLAST

GPU-HMMER

mCUDA-MEME

SeqNFind

UGENE

WideLM

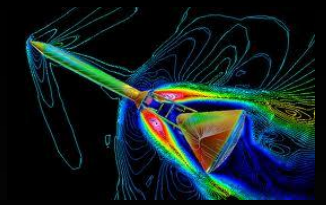

Computational Fluid Dynamics

ANSYS Fluent

Altair AcuSolve

FluiDyna Culises - OpenFOAM

Turbostream

Prometech Particleworks

Computational Structural Mechanics

ANSYS Mechanical

Abaqus/Standard

MSC NASTRAN

Altair – RADIOSS

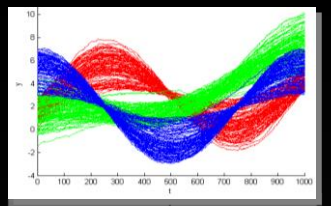

Numerical Analysis

MATLAB

Jacket for MATLAB

Mathematica

Physics

CHROMA

GTC

MILC

Climate/Weather/ Ocean Modeling

COSMO

WRF

NEMO

ASUCA (JMA

NIM

HOMME

CAM/SE

GEOS-5

Download Catalog of GPU Apps

Resource

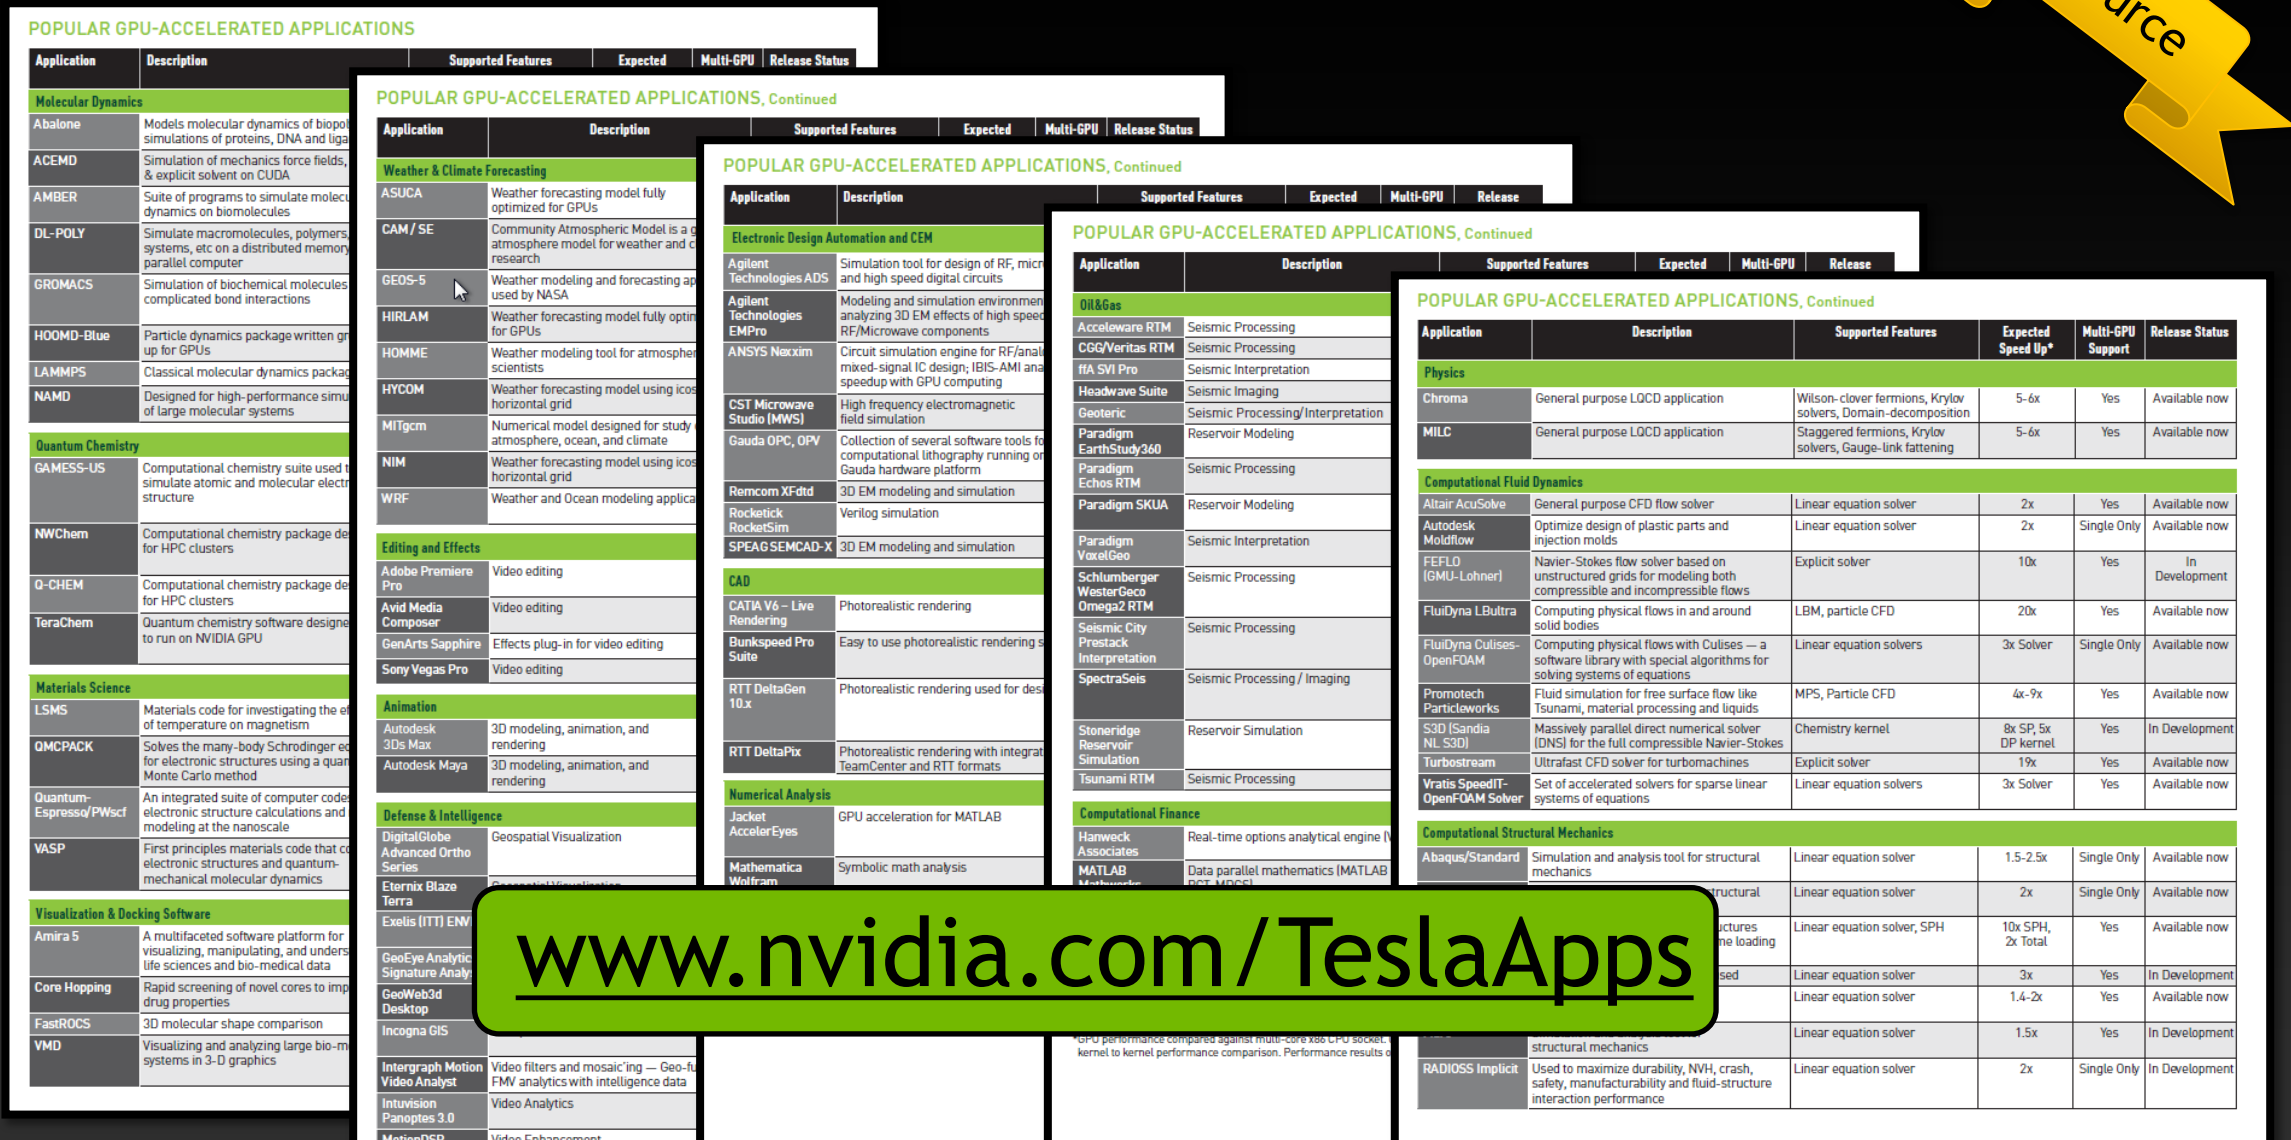

POPULAR GPU-ACCELERATED APPLICATIONS

Application	Description	Supported Features	Expected	Multi-GPU	Release Status
Molecular Dynamics					
Abalone	Models molecular dynamics of biopol simulations of proteins, DNA and liga				
ACEMD	Simulation of mechanics force fields, & explicit solvent on CUDA				
AMBER	Suite of programs to simulate molecu dynamics on biomolecules				
DL-POLY	Simulate macromolecules, polymers, systems, etc on a distributed memory parallel computer				
GROMACS	Simulation of biochemical molecules complicated bond interactions				
HOOMD-Blue	Particle dynamics package written gr up for GPUs				
LAMMPS	Classical molecular dynamics packag				
NAMD	Designed for high-performance simu of large molecular systems				
Quantum Chemistry					
GAMESS-US	Computational chemistry suite used t simulate atomic and molecular electr structure				
NWChem	Computational chemistry package de for HPC clusters				
Q-CHEM	Computational chemistry package de for HPC clusters				
TeraChem	Quantum chemistry software designe to run on NVIDIA GPU				
Materials Science					
LSMS	Materials code for investigating the ef of temperature on magnetism				
QMCPACK	Solves the many-body Schrodinger eq for electronic structures using a quan Monte Carlo method				
Quantum-Espresso/PWscf	An integrated suite of computer code electronic structure calculations and modeling at the nanoscale				
VASP	First principles materials code that c electronic structures and quantum-mechanical molecular dynamics				
Visualization & Docking Software					
Amira 5	A multifaceted software platform for visualizing, manipulating, and unders life sciences and bio-medical data				
Core Hopping	Rapid screening of novel cores to imp drug properties				
FastROCS	3D molecular shape comparison				
VMD	Visualizing and analyzing large bio-m systems in 3-D graphics				

POPULAR GPU-ACCELERATED APPLICATIONS, Continued

Application	Description	Supported Features	Expected	Multi-GPU	Release Status
Weather & Climate Forecasting					
ASUCA	Weather forecasting model fully optimized for GPUs				
CAM / SE	Community Atmospheric Model is a g atmosphere model for weather and c research				
GEOS-5	Weather modeling and forecasting ap used by NASA				
HIRLAM	Weather forecasting model fully optim for GPUs				
HOMME	Weather modeling tool for atmospher scientists				
HYCOM	Weather forecasting model using icos horizontal grid				
MITgcm	Numerical model designed for study atmosphere, ocean, and climate				
NIM	Weather forecasting model using icos horizontal grid				
WRF	Weather and Ocean modeling applica				
Editing and Effects					
Adobe Premiere Pro	Video editing				
Avid Media Composer	Video editing				
GenArts Sapphire	Effects plug-in for video editing				
Sony Vegas Pro	Video editing				
Animation					
Autodesk 3Ds Max	3D modeling, animation, and rendering				
Autodesk Maya	3D modeling, animation, and rendering				
Defense & Intelligence					
DigitalGlobe Advanced Ortho Series	Geospatial Visualization				
Eternix Blaze Terra					
Exelis (ITT) ENV					
GeoEye Analytic Signature Analy					
GeoWeb3d Desktop					
Incogna GIS					
Intergraph Motion Video Analyst	Video filters and mosaic'ing — Geo-fu FMV analytics with intelligence data				
Intuvision Panoptes 3.0	Video Analytics				
MotionDSP	Video Enhancement				

POPULAR GPU-ACCELERATED APPLICATIONS, Continued

Application	Description	Supported Features	Expected	Multi-GPU	Release
Electronic Design Automation and CEM					
Agilent Technologies ADS	Simulation tool for design of RF, micr and high speed digital circuits				
Agilent Technologies EMPro	Modeling and simulation environmen analyzing 3D EM effects of high speed RF/Microwave components				
ANSYS Nexxim	Circuit simulation engine for RF/anal mixed-signal IC design; IBIS-AMI ana speedup with GPU computing				
CST Microwave Studio (MWS)	High frequency electromagnetic field simulation				
Gauda OPC, OPV	Collection of several software tools fo computational lithography running or Gauda hardware platform				
Remcom XFdtd	3D EM modeling and simulation				
Rocketick RocketSim	Verilog simulation				
SPEAG SEMCAD-X	3D EM modeling and simulation				
CAD					
CATIA V6 – Live Rendering	Photorealistic rendering				
Bunkspeed Pro Suite	Easy to use photorealistic rendering s				
RTT DeltaGen 10.x	Photorealistic rendering used for des				
RTT DeltaPix	Photorealistic rendering with integrat TeamCenter and RTT formats				
Numerical Analysis					
Jacket AccelerEyes	GPU acceleration for MATLAB				
Mathematica Wolfram	Symbolic math analysis				

POPULAR GPU-ACCELERATED APPLICATIONS, Continued

Application	Description	Supported Features	Expected	Multi-GPU	Release
Oil&Gas					
Acceleware RTM	Seismic Processing				
CGG/Veritas RTM	Seismic Processing				
ffA SVI Pro	Seismic Interpretation				
Headwave Suite	Seismic Imaging				
Geoteric	Seismic Processing/Interpretation				
Paradigm EarthStudy360	Reservoir Modeling				
Paradigm Echos RTM	Seismic Processing				
Paradigm SKUA	Reservoir Modeling				
Paradigm VoxelGeo	Seismic Interpretation				
Schlumberger WesterGeco Omega2 RTM	Seismic Processing				
Seismic City Prestack Interpretation	Seismic Processing				
SpectraSeis	Seismic Processing / Imaging				
Stoneridge Reservoir Simulation	Reservoir Simulation				
Tsunami RTM	Seismic Processing				
Computational Finance					
Hanweck Associates	Real-time options analytical engine [				
MATLAB	Data parallel mathematics (MATLAB				

*GPU performance compared against multi-core x86 CPU socket. kernel to kernel performance comparison. Performance results o

POPULAR GPU-ACCELERATED APPLICATIONS, Continued

Application	Description	Supported Features	Expected Speed Up*	Multi-GPU Support	Release Status
Physics					
Chroma	General purpose LQCD application	Wilson-clover fermions, Krylov solvers, Domain-decomposition	5-6x	Yes	Available now
MILC	General purpose LQCD application	Staggered fermions, Krylov solvers, Gauge-link fattening	5-6x	Yes	Available now
Computational Fluid Dynamics					
Altair AcuSolve	General purpose CFD flow solver	Linear equation solver	2x	Yes	Available now
Autodesk Moldflow	Optimize design of plastic parts and injection molds	Linear equation solver	2x	Single Only	Available now
FEFLO (GMU-Lohner)	Navier-Stokes flow solver based on unstructured grids for modeling both compressible and incompressible flows	Explicit solver	10x	Yes	In Development
FluiDyna LBultra	Computing physical flows in and around solid bodies	LBM, particle CFD	20x	Yes	Available now
FluiDyna Culises-OpenFOAM	Computing physical flows with Culises — a software library with special algorithms for solving systems of equations	Linear equation solvers	3x Solver	Single Only	Available now
Promotech Particleworks	Fluid simulation for free surface flow like Tsunami, material processing and liquids	MPS, Particle CFD	4x-9x	Yes	Available now
S3D (Sandia NL S3D)	Massively parallel direct numerical solver (DNS) for the full compressible Navier-Stokes	Chemistry kernel	8x SP, 5x DP kernel	Yes	In Development
Turbostream	Ultrafast CFD solver for turbomachines	Explicit solver	19x	Yes	Available now
Vratis SpeedIT-OpenFOAM Solver	Set of accelerated solvers for sparse linear systems of equations	Linear equation solvers	3x Solver	Yes	Available now
Computational Structural Mechanics					
Abaqus/Standard	Simulation and analysis tool for structural mechanics	Linear equation solver	1.5-2.5x	Single Only	Available now
	structural	Linear equation solver	2x	Single Only	Available now
	uctures me loading	Linear equation solver, SPH	10x SPH, 2x Total	Yes	Available now
	sed	Linear equation solver	3x	Yes	In Development
		Linear equation solver	1.4-2x	Yes	Available now
	structural mechanics	Linear equation solver	1.5x	Yes	In Development
RADIOSS Implicit	Used to maximize durability, NVH, crash, safety, manufacturability and fluid-structure interaction performance	Linear equation solver	2x	Single Only	In Development

www.nvidia.com/TeslaApps

Steady March of the GPU Performance with Apps

96 server nodes
46 ns/day

2009

8 Tesla M2050
52 ns/day

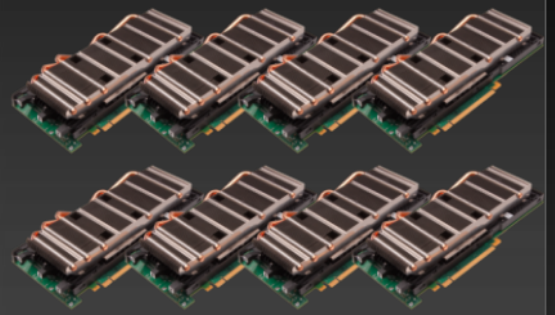

2010

4 Tesla M2090
69 ns/day

2011

2 Tesla K10 GPUs
66 ns/day

2012

AMBER(beta)- JAC NVE Benchmark

AMBER Benchmark Report
Accelerates Simulations of All Sizes

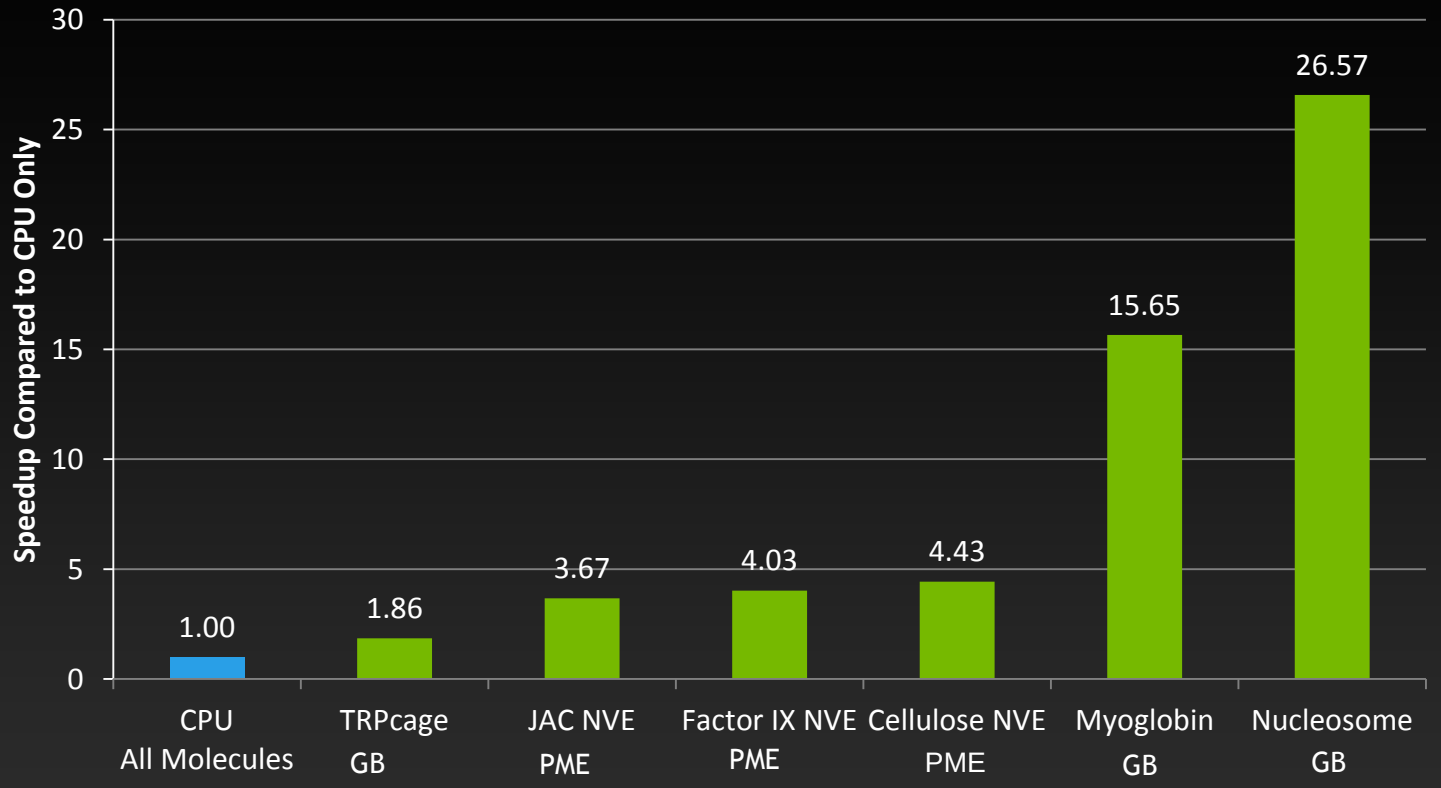

Running AMBER 12 with CUDA 4.2 ECC Off

The blue node contains 2x Intel X5670 CPUs (6 Cores per CPU)

Each green node contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 1x NVIDIA K10 GPU

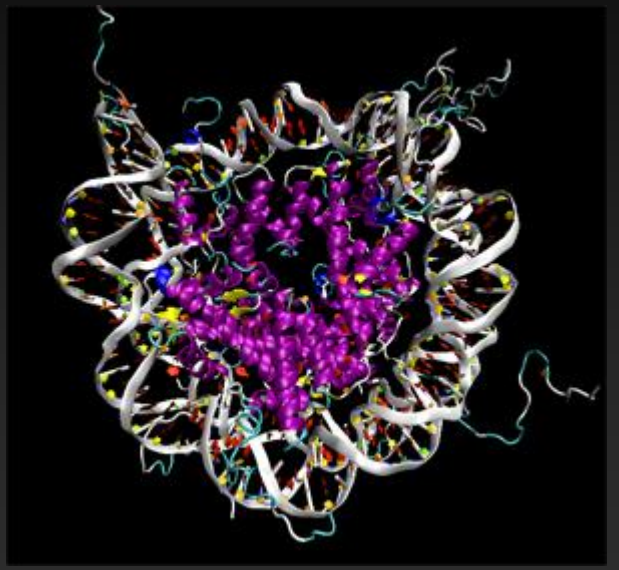

Nucleosome

Gain 26x performance by adding just 1 K10 GPU
compared to dual CPU performance

AMBER Benchmark Report
Strong Scaling over Nodes

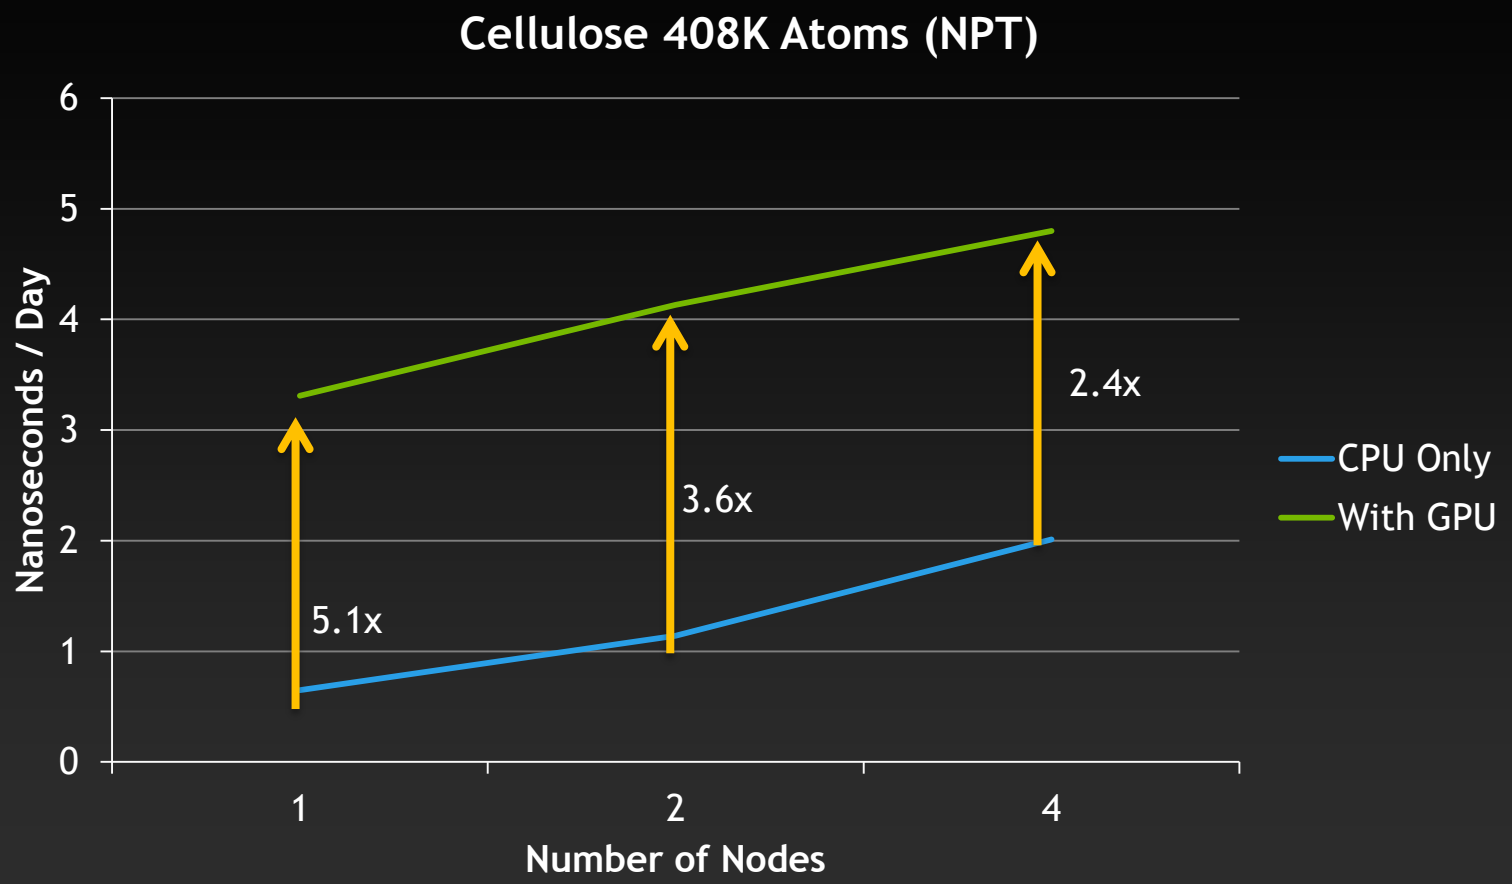

Running AMBER 12 with CUDA 4.2 ECC Off

The blue nodes contains 2x Intel X5670 CPUs (6 Cores per CPU)

The green nodes contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 2x NVIDIA K10 GPUs

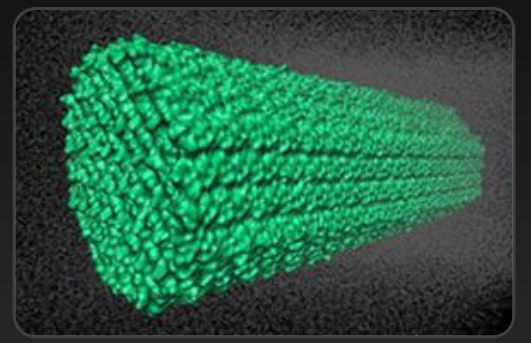

Cellulose

GPUs significantly outperform CPUs while scaling over multiple nodes

AMBER Benchmark Report
Replace 5 Nodes with 1 GPU

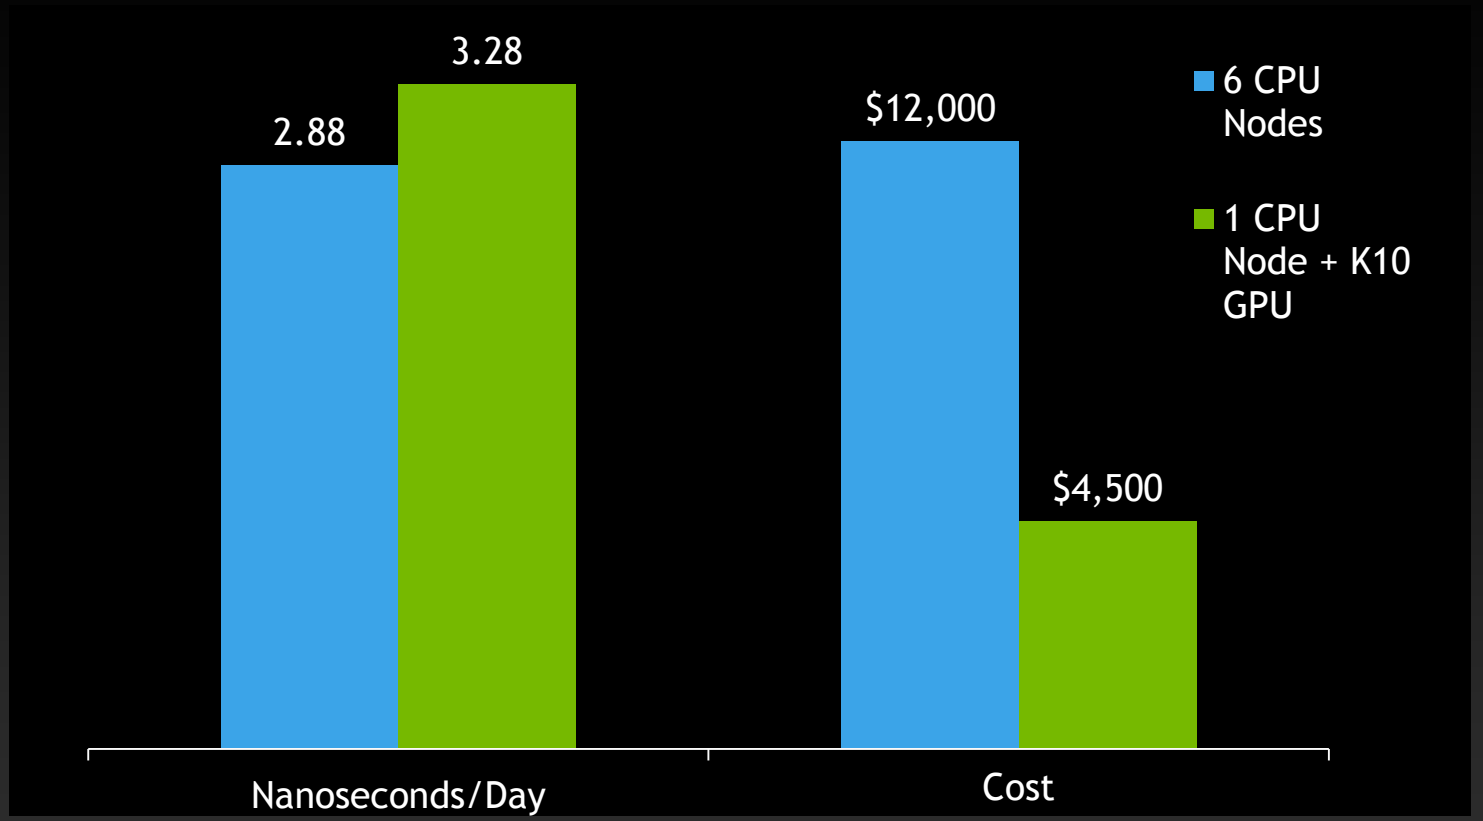

Running AMBER 12 with CUDA 4.2 ECC Off

Each blue node contains 2x Intel X5670 CPUs (6 Cores, $1000 per CPU)

The green node contains 2x Intel X5670 CPUs (6 Cores, $1000 per CPU) and 1x NVIDIA K10 GPU ($2500 per GPU)

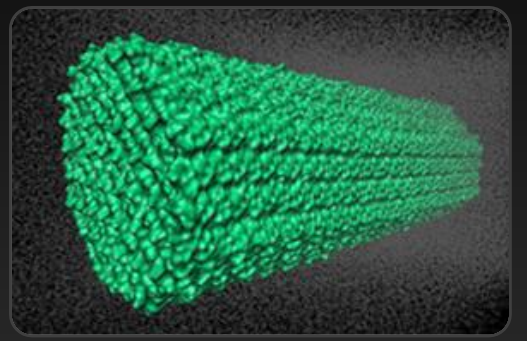

Cellulose

Cut down simulation costs to ⅓ and gain higher performance

AMBER Benchmark Report
Greener Science

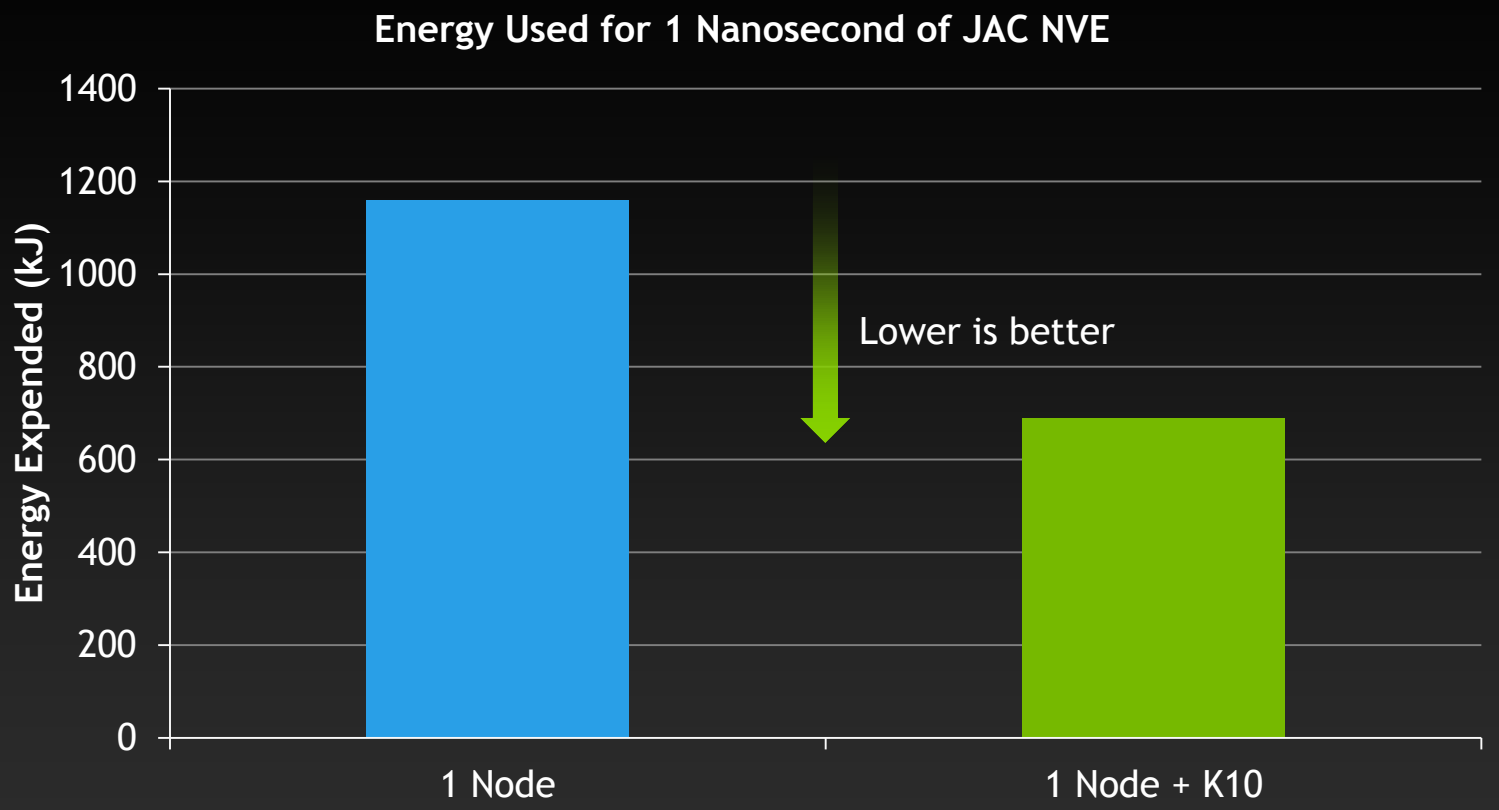

Running AMBER 12 with CUDA 4.2 ECC Off

The blue node contains 2x Intel X5670 CPUs (95W each, 6 Cores per CPU)

The green node contains 2x Intel X5670 CPUs (95W each, 6 Cores per CPU) and 1x NVIDIA K10 GPU (225W per GPU)

Energy Expended = Power x Time

The GPU Accelerated system uses 40% less energy

Download GPU Apps Performance Reports

Resource

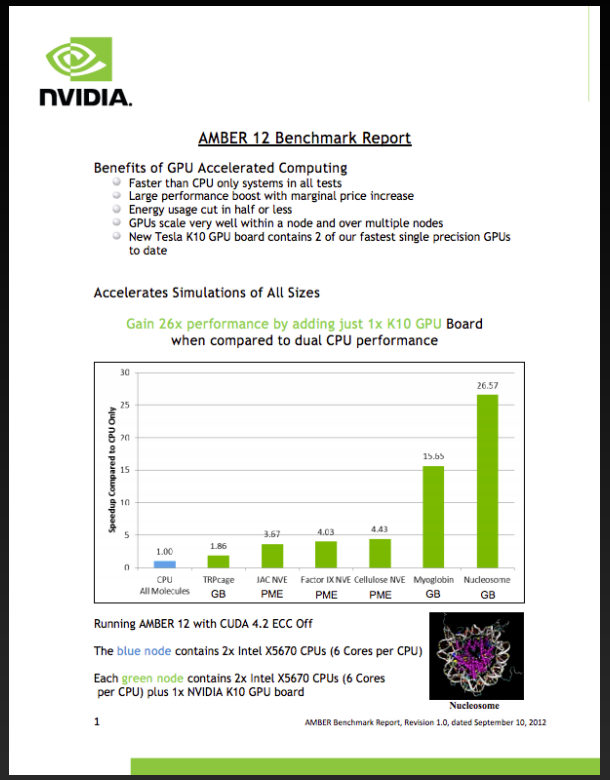

NVIDIA.

AMBER 12 Benchmark Report

Benefits of GPU Accelerated Computing

- Faster than CPU only systems in all tests
- Large performance boost with marginal price increase
- Energy usage cut in half or less
- GPUs scale very well within a node and over multiple nodes
- New Tesla K10 GPU board contains 2 of our fastest single precision GPUs to date

Accelerates Simulations of All Sizes

Gain 26x performance by adding just 1x K10 GPU Board when compared to dual CPU performance

Running AMBER 12 with CUDA 4.2 ECC Off

The blue node contains 2x Intel X5670 CPUs (6 Cores per CPU)

Each green node contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 1x NVIDIA K10 GPU board

Nucleosome

1

AMBER Benchmark Report, Revision 1.0, dated September 10, 2012

AMBER
NAMD
GROMACS
LAMMPS

More coming soon…

www.nvidia.com/GPUTestDrive

Impact on Academic Labs

> “Students are now able to run calculations that would not have been feasible on our local CPU-based resources before. Research throughput the group has been enhanced significantly.
> Simply a great investment!”

Jodi Hadden
Chemistry Graduate Student
Woods Computing Lab
Complex Carbohydrate Research Center
University of Georgia

The University of Georgia
1785

Preventing New H1N1 Epidemics

- Drugs became ineffective due to rapid H1N1 mutations
- Without access to a large cluster, researchers turned to GPUs to decode molecular structure of H1N1 virus mutations
- GPUs shortened research time by 50% using a small cluster with 8 GPUs
- Discovered new ways to design inhibitor drugs to quickly address mutations

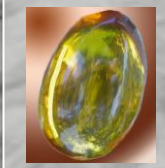

AMBER (PMEMD)

University of BRISTOL

Accelerating Diagnosis of Cancer Cells using MATLAB with GPUs

- Vector quantization – one of the diagnostics techniques for breast cancer diagnosis that requires processing high resolution images
- 230,480 women were diagnosed with breast cancer in 2011 in US and millions of images need to be processed for diagnosis each year
- Researchers from Cornell, Univ. of Michigan and Rutgers Univ. using MATLAB with GPUs have been able to reduce processing time from 86.9 to 5.9 seconds – 14 times faster

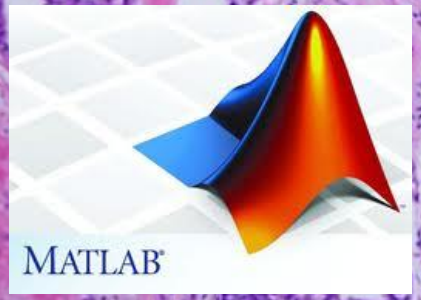

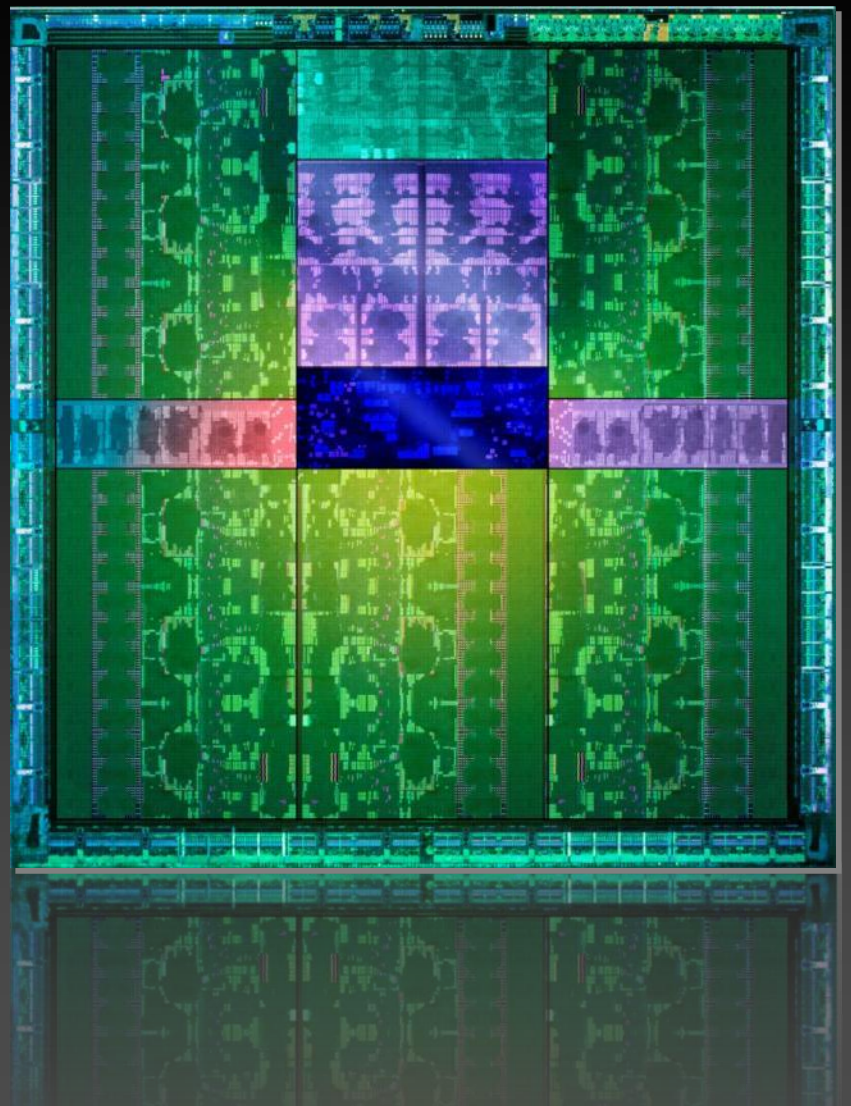

KEPLER

THE WORLD'S FASTEST, MOST EFFICIENT HPC ACCELERATOR

3x Perf/Watt

Max Utilization

Easier to Program

Tesla Data Center & Workstation GPU Solutions

Servers & Blades

		K10	M2090	M2075
Cores		3072	512	448
Memory		8 GB	6 GB	6 GB
Memory bandwidth (ECC off)		320 GB/s	177.6 GB/s	150 GB/s
Peak Perf Gflops	Single Precision	4580	1331	1030
	Double Precision	190	665	515

Workstations

C2075
448
6 GB
148.8 GB/s
1030
515

Tesla K10

3x Single Precision

1.8x Memory Bandwidth

Image, Signal, Seismic

Available Now

Tesla K20

3x Double Precision

Hyper-Q, Dynamic Parallelism

CFD, FEA, Finance, Physics

Available Soon
Pre-order now

Test Drive GPU Apps!

Experience The Acceleration

Resource

Sign up for FREE GPU Test Drive on remotely hosted clusters

Pre-loaded Applications:

AMBER, NAMD, LAMMPS, Quantum Espresso, Terachem…

"We are very happy with the performance of the Fermi cards, and the additional speed provided by Kepler is definitely welcome"

Lars Pastewka, LAMMPS user from Johns Hopkins University

www.nvidia.com/GPUTestDrive

Contact:
Devang Sachdev
dsachdev@nvidia.com

Thanks

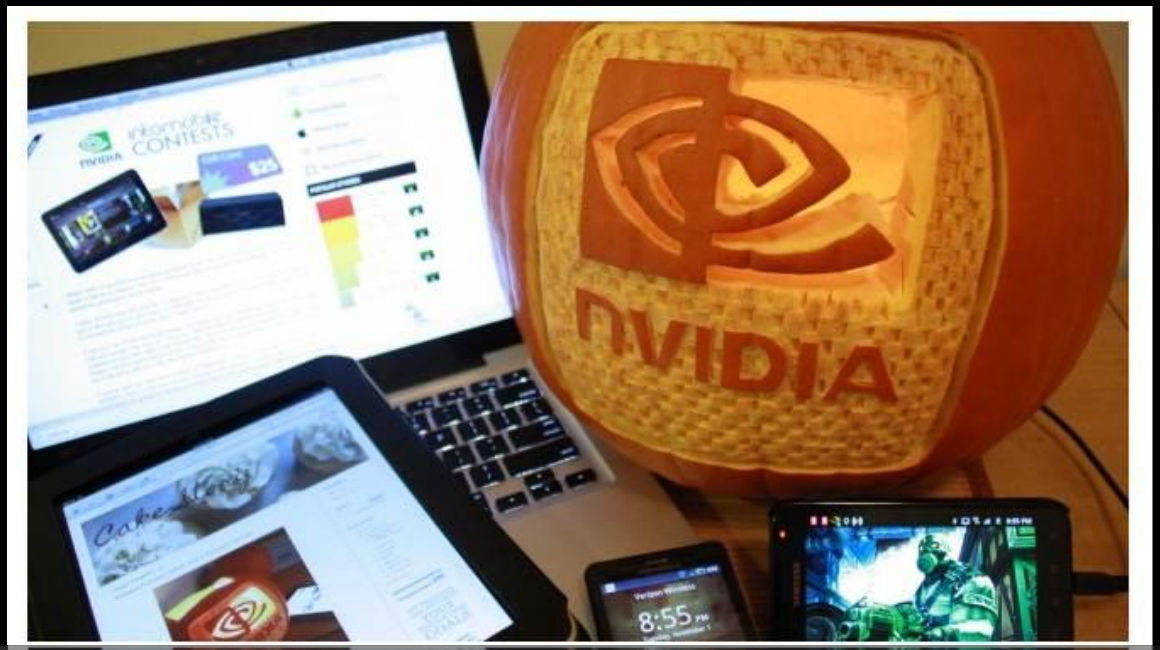

Hear from Luminaries and Experts on Accelerated Computing

GPU Technology Theater

November 12-15, 2012
NVIDIA Booth #2217 @ SC12
Salt Lake City, Utah

Science and developer talks every 30 minutes

Live webcast and interactive Q&A available!

Schedule and more details:
www.nvidia.com/sc12

SC12
Salt Lake City, Utah

GPU Technology Conference

March 18 - 21, 2013
San Jose Convention Center

The smartest people
The brightest ideas
The biggest opportunities

Schedule and registration coming soon:
www.gputechconf.com

GPU TECHNOLOGY CONFERENCE

Backup

Summary: H1N1 Bristol University

Problem

- 2009 outbreak of H1N1 influenza virus infected killed up to 18,300 people
- Virus mutations rendered widely used life saving drugs Tamiflu and Relenza ineffective
- Difficult to study virus/mutations in wet labs, simulation is the only option
- Struggling to make progress as simulations require access to large computing resources
- Researchers from UK and Thai universities needed full time access to entire CPU cluster for 2 months for simulations
- Monopolizing computing resources shared across the lab not feasible

Solution

- With GPUs researchers could run, repeat much larger number of simulations using one-fifth the servers in one-half the time (1 month vs. 2)

Benefit

- Identify new ways in which inhibitor drugs can be quickly designed to address these mutations
- With ~1K fatalities/month related to H1N1 at outbreak, reduction in drug discovery even by a month can make a huge impact on prevention of future epidemics
- Small, affordable GPU-based servers give researchers dedicated access to a high-performance systems in-house to help them make discoveries faster and potentially prevent epidemics

Key Application Developers Embracing GPUs

"(A petascale GPU cluster) … would revolutionize the approach taken for studying enzyme drug binding, materials science and other fields such as enzyme catalysis for things such as biofuel work."

Dr. Ross Walker, Adjunct Professor of Chemistry, UCSD
Lead developer for AMBER

"Our current GPU code on C2050 is equivalent to close to 100 cores on NICS Kraken."

Dr. Yifent Cui, Lead for High Performance Geocomputing Lab
Lead developer for AWP-ODC

"We plan to outfit GADGET3 with GPU support for all of its 3 primary computational intensive parts."

Dr. Volker Springel,
Professor at Heidelberg Instituite of Theoretical Studies &
Lead developer for GADGET

Explosive Growth of GPU Computing

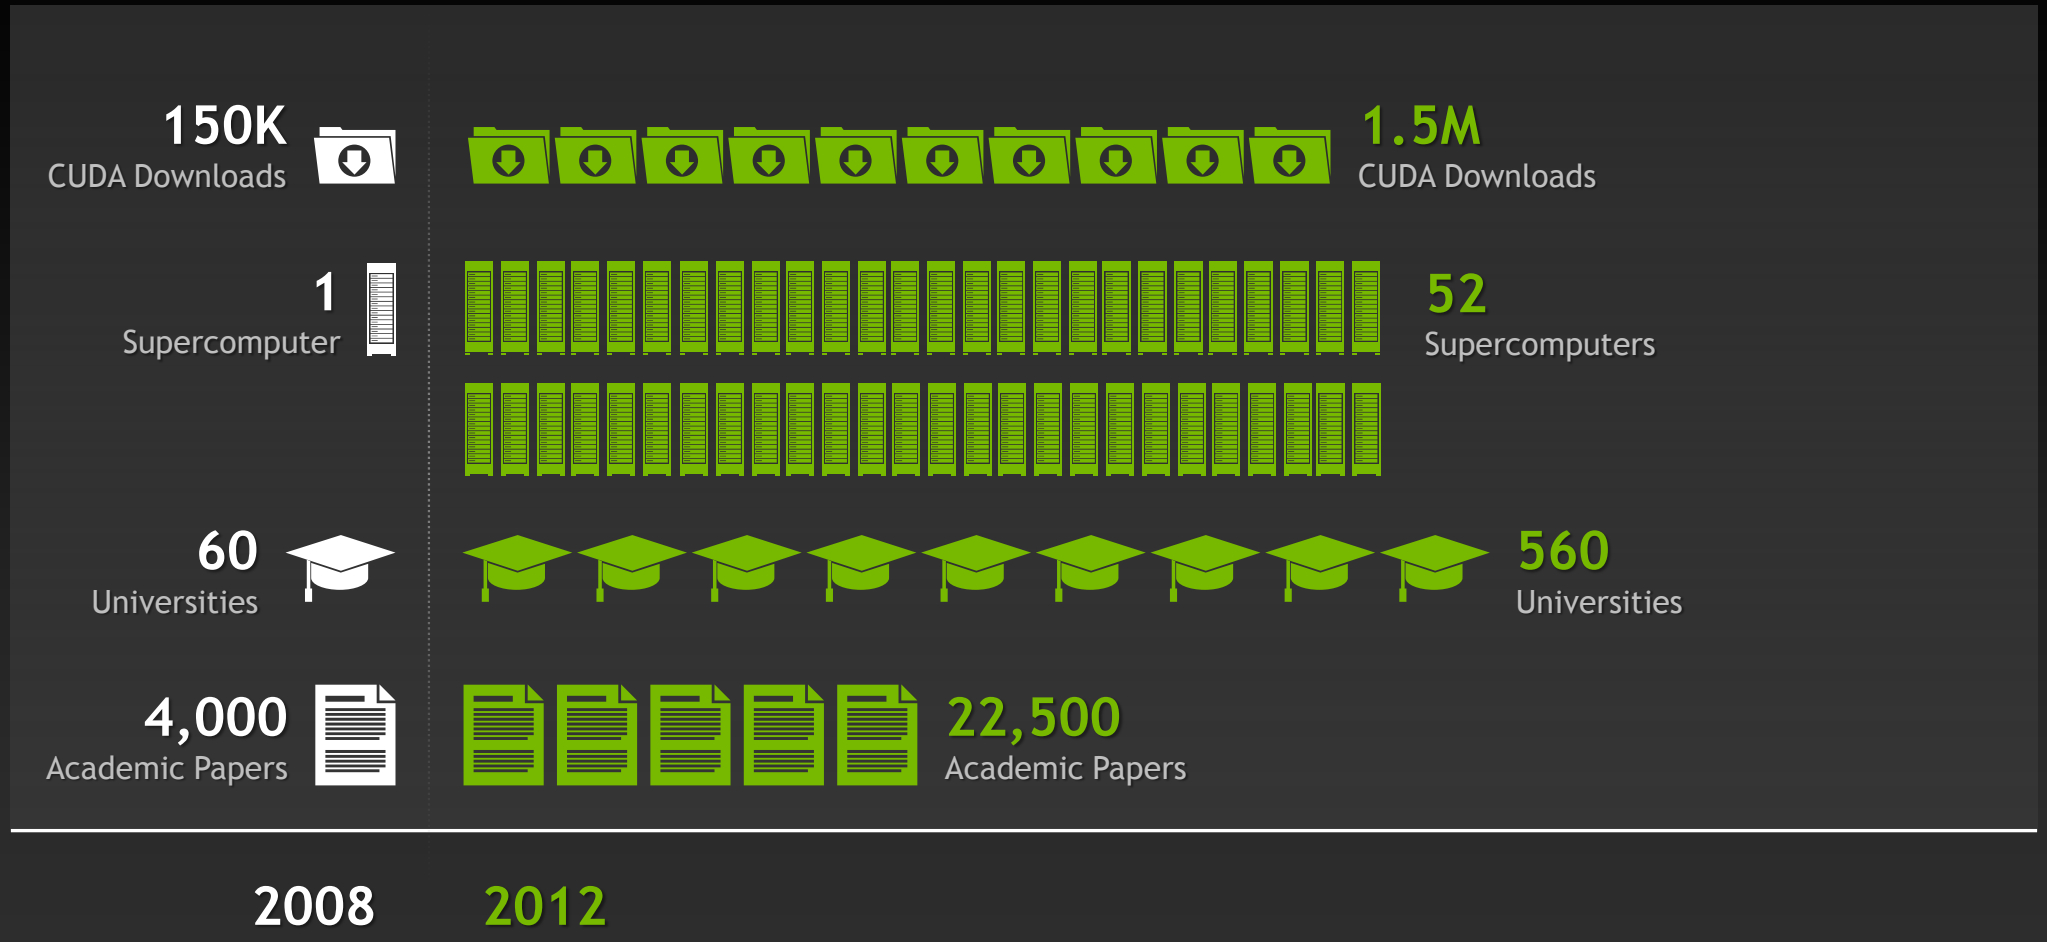

AMBER Benchmark Report
Extreme Performance

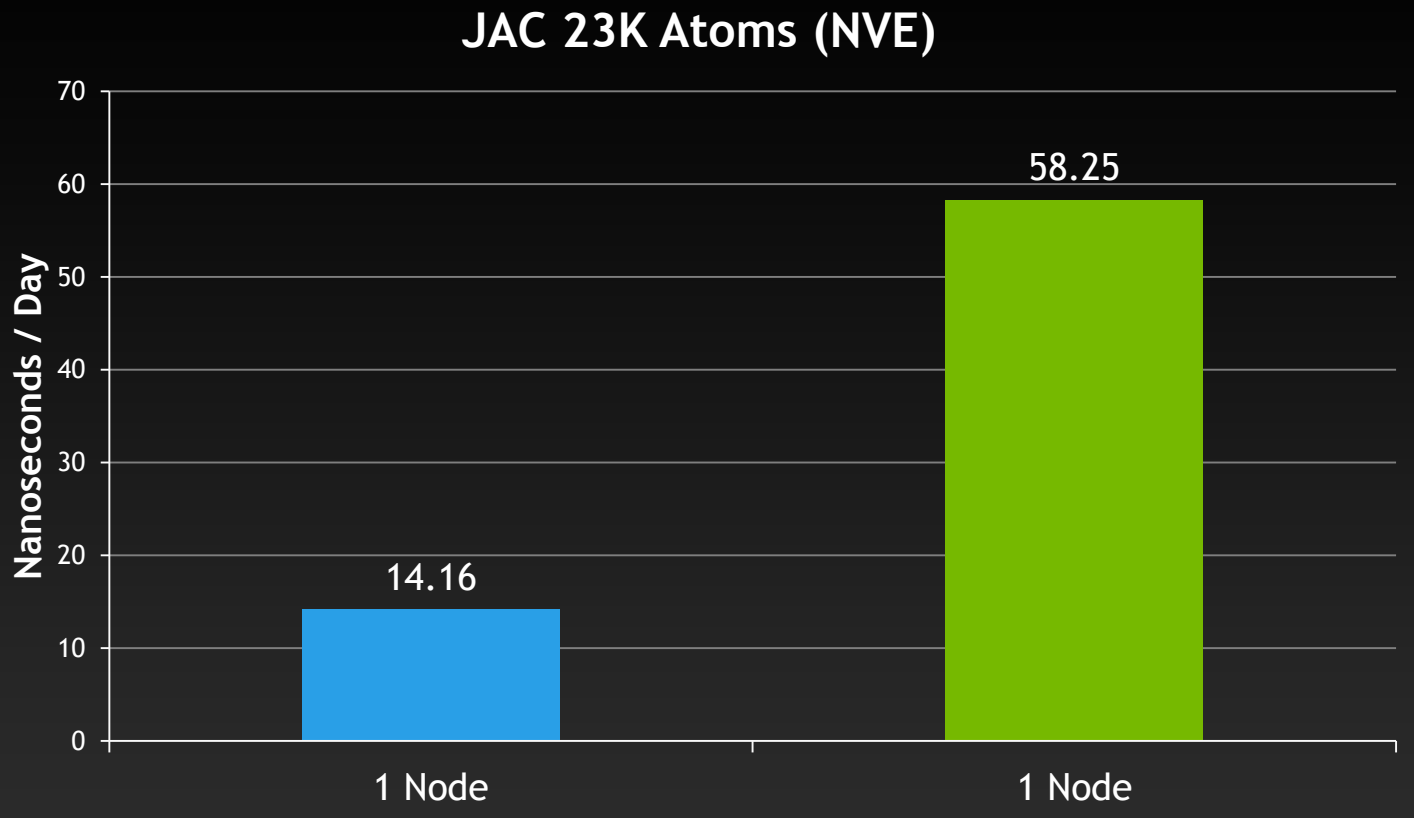

Running AMBER 12 with CUDA 4.2 ECC Off

The blue node contains 2x Intel X5670 CPUs (6 Cores per CPU)

The green node contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 2x NVIDIA K10 GPUs

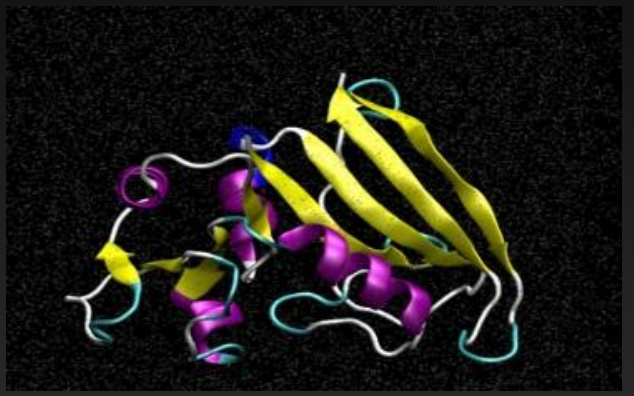

DHFR

Gain 4X performance by adding just 2 GPUs
when compared to dual CPU performance

Quick Poll

Do you write your own HPC applications?

- Yes
- No – I run HPC applications available for download/license from ISVs