6th September 2011 – Belgrade

University of Belgrade, Mathematics faculty

NVIDIA HARDWARE FOR HIGH PERFORMANCE COMPUTING

Piero Altoè – HPC Sales Manager (EMEA)

TESLA

GPU Computing

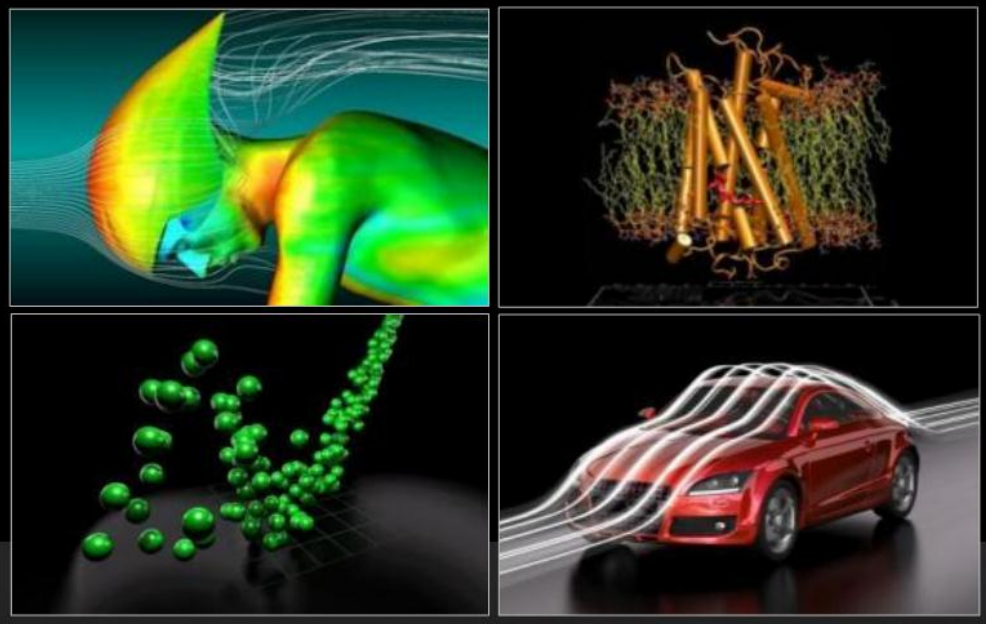

Accelerating High Performance Computing

http://www.nvidia.com/tesla

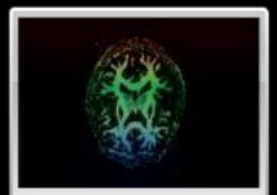

146X

Medical Imaging
U of Utah

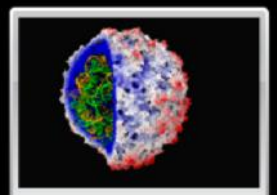

36X

Molecular Dynamics
U of Illinois, Urbana

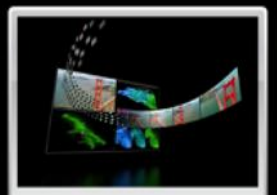

18X

Video Transcoding
Elemental Tech

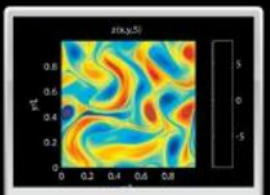

50X

Matlab Computing
AccelerEyes

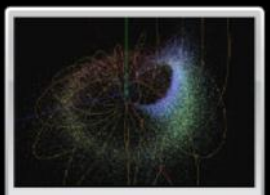

100X

Astrophysics
RIKEN

GPUs Accelerate Science

149X

Financial Simulation
Oxford

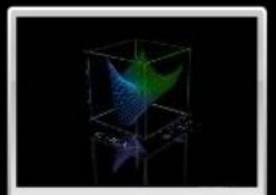

47X

Linear Algebra
Universidad Jaime

20X

3D Ultrasound
Techniscan

130X

Quantum Chemistry
U of Illinois, Urbana

30X

Gene Sequencing
U of Maryland

Minimum Change, Big Speed-up

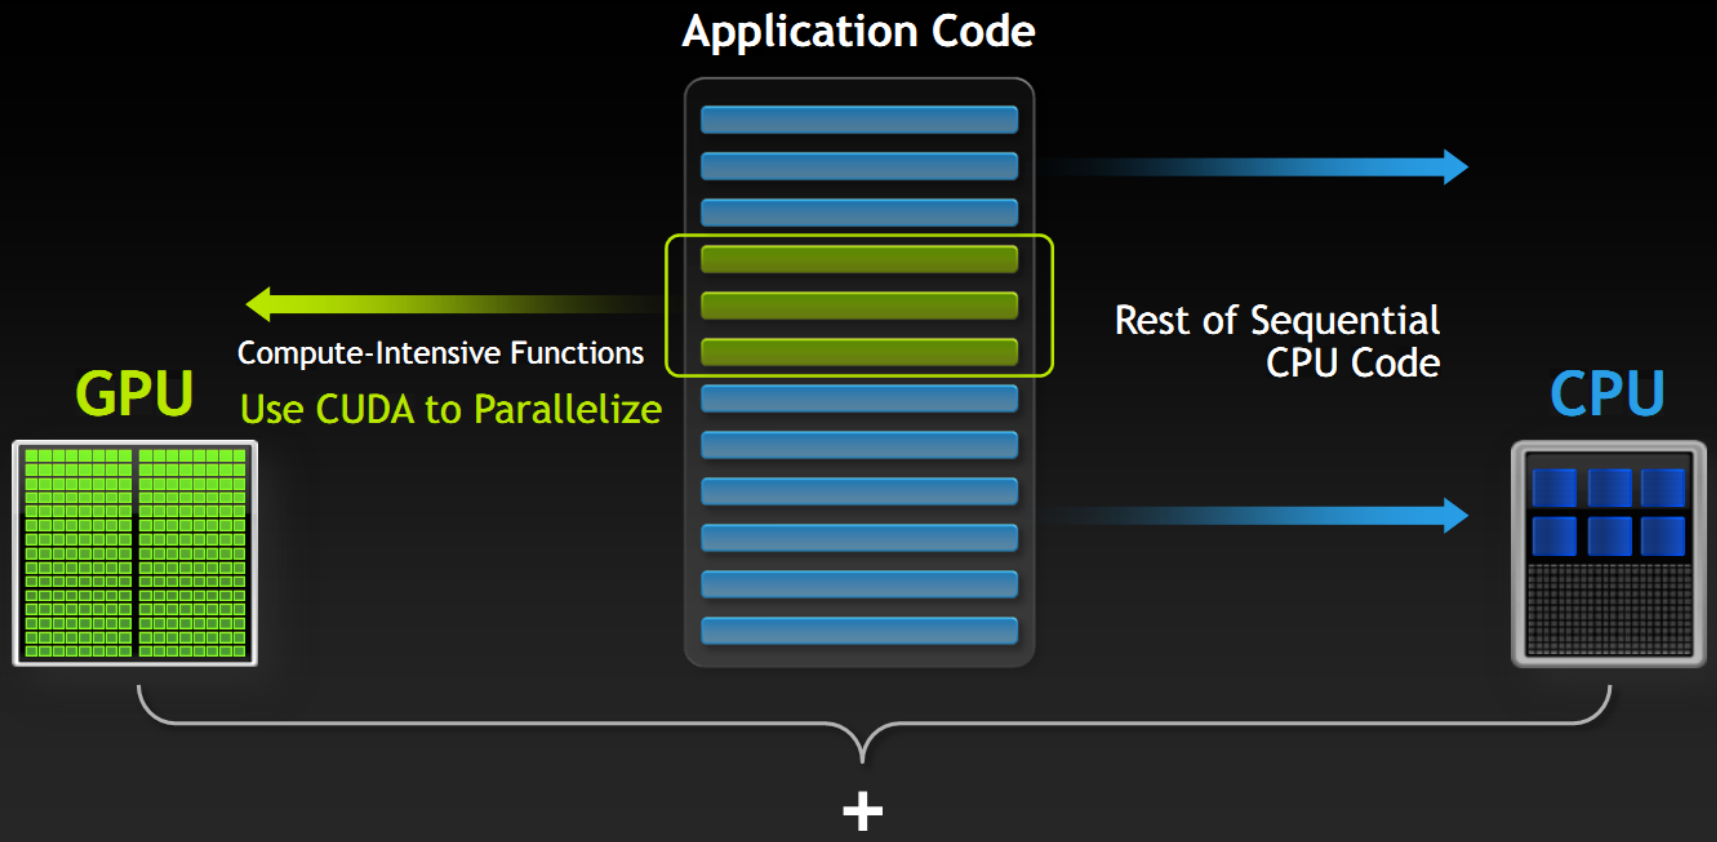

Tesla GPUs Power 3 of Top 5 Supercomputers

#1 : Tianhe-1A
7168 Tesla GPU's 2.5 PFLOPS

#3 : Nebulae
4650 Tesla GPU's 1.2 PFLOPS

#4 : Tsubame 2.0
4224 Tesla GPU's 1.194 PFLOPS

"We not only created the world's fastest computer, but also implemented a heterogeneous computing architecture incorporating CPU and GPU, this is a new innovation."

Premier Wen Jiabao
Public comments acknowledging Tianhe-1A

World's Fastest MD Simulation

Sustained Performance of 1.87 Petaflops/s

Institute of Process Engineering (IPE)

Chinese Academy of Sciences (CAS)

MD Simulation for Crystalline Silicon

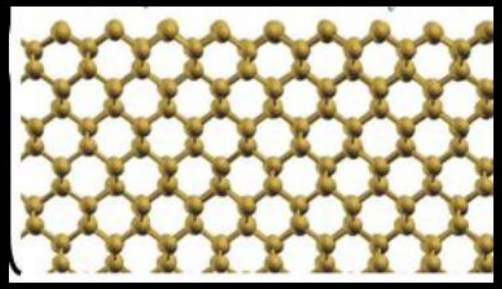

Used all 7168 Tesla GPUs on Tianhe-1A GPU Supercomputer

World's Greenest Petaflop Supercomputer

Tsubame 2.0

Tokyo Institute of Technology

- 1.19 Petaflops
- 4,224 Tesla M2050 GPUs

Commercial Apps Accelerated by GPUs

Molecular Dynamics	AMBER CHARMM DL_POLY GROMACS LAMMPS NAMD
Fluid Dynamics	Altair Acusolve Autodesk Moldflow OpenFOAM Prometech Particlework Turbostream
Earth Sciences	ASUCA HOMME NASA GEOS-5 NOAA NIM WRF
Engineering Simulation	Agilent EMPro ANSYS Mechanical ANSYS Nexxim CST Microwave Studio Impetus AFEA Remcom XFdtd SIMULIA Abaqus
Others	GADGET2 MATLAB Mathematica NBODY Paradigm VoxelGeo PARATEC Schlumberger Petrel

Widespread Adoption of GPUs

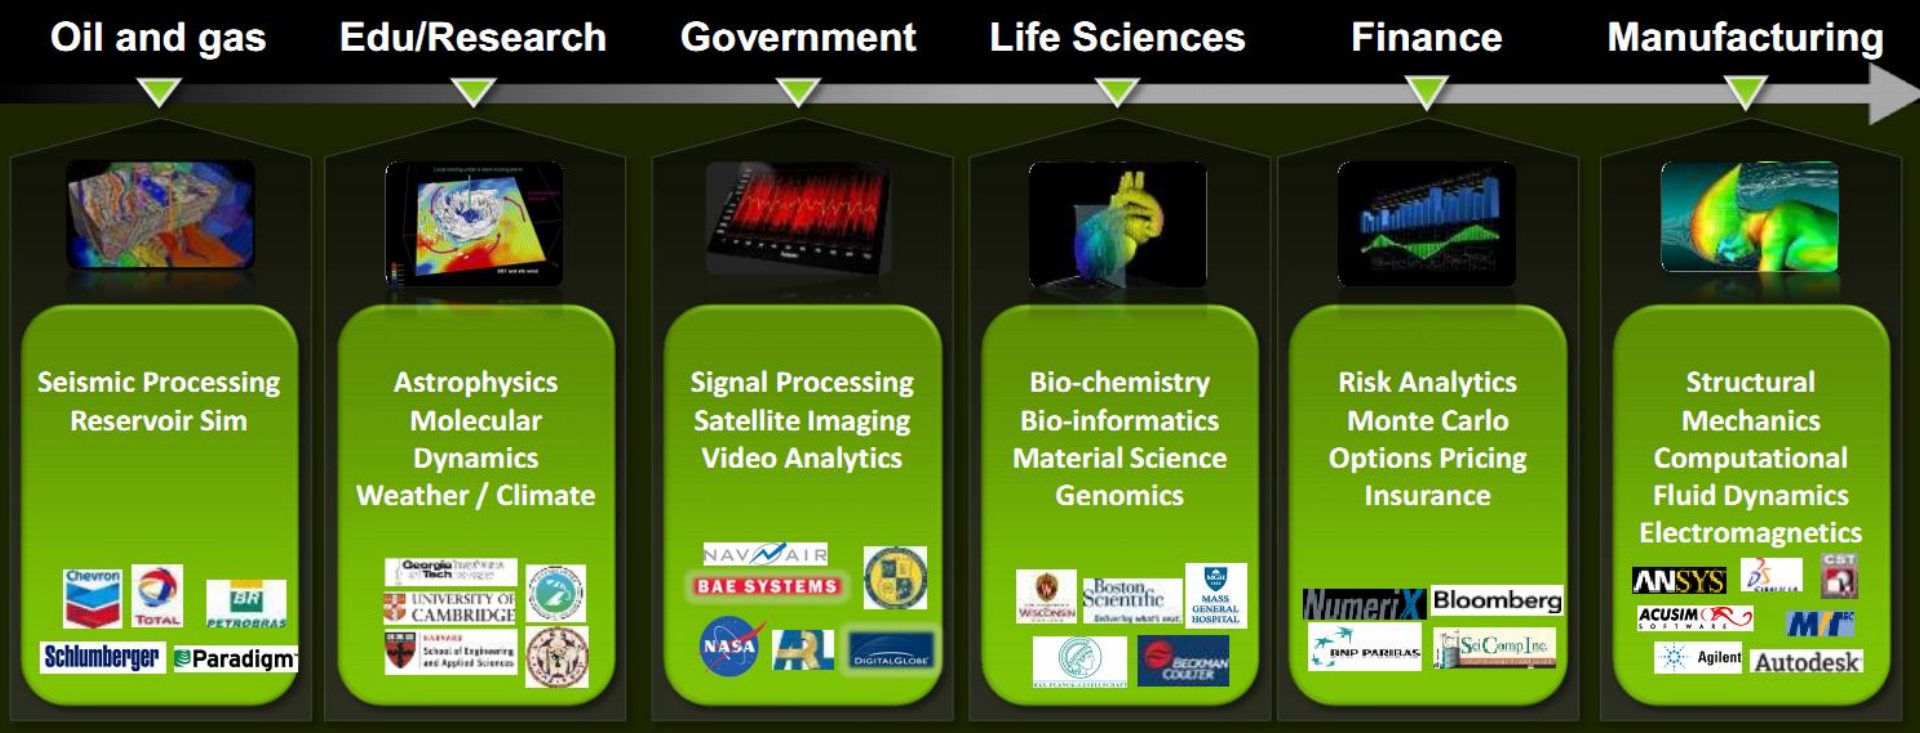

CUDA: Easy to Use Parallel Programming Model

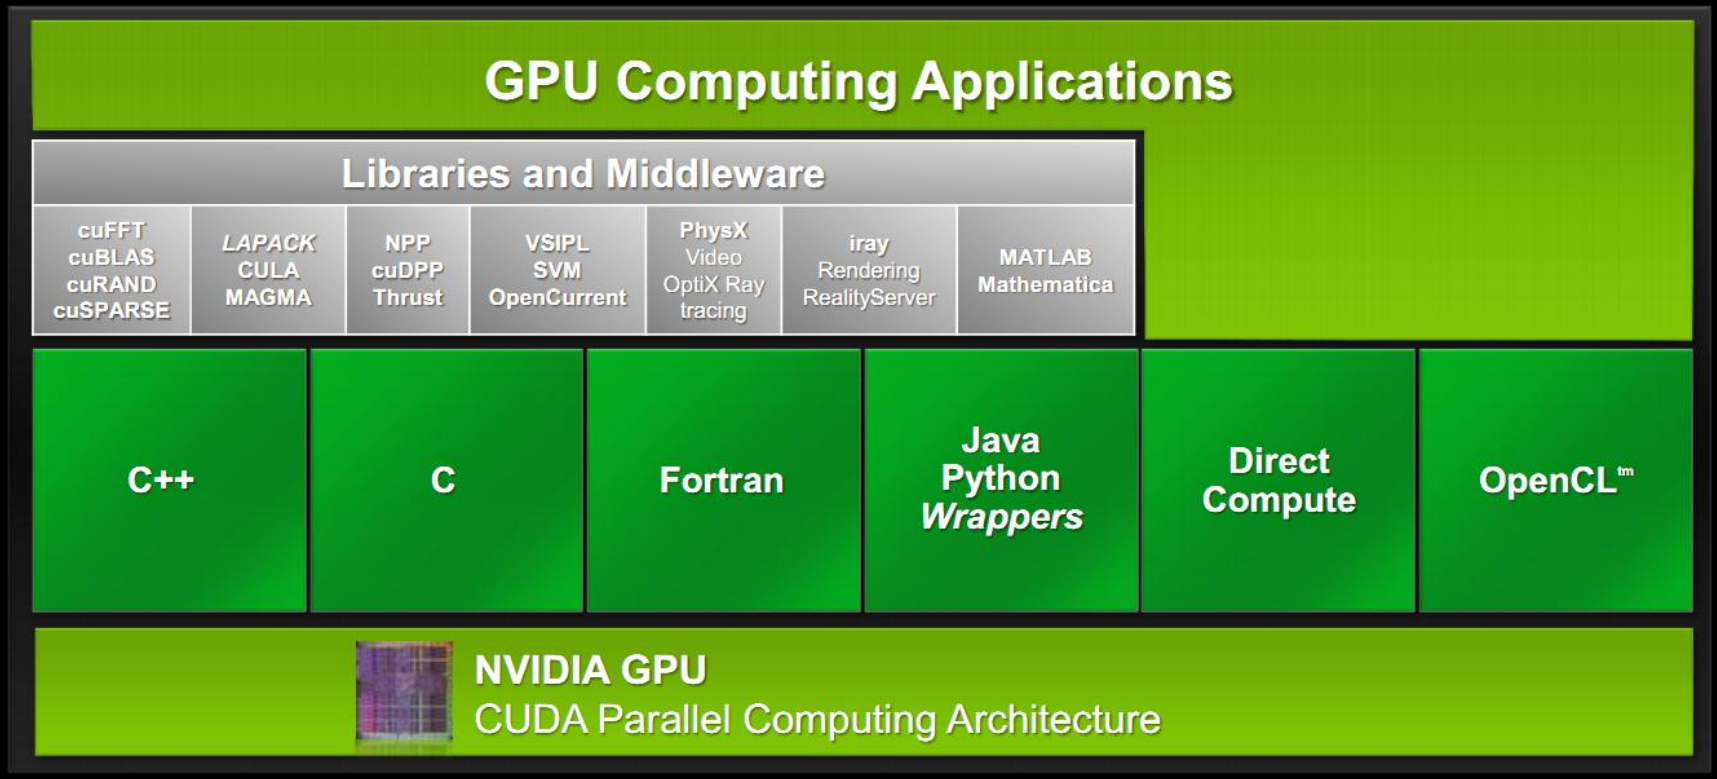

Increasing Number of Professional CUDA Applications

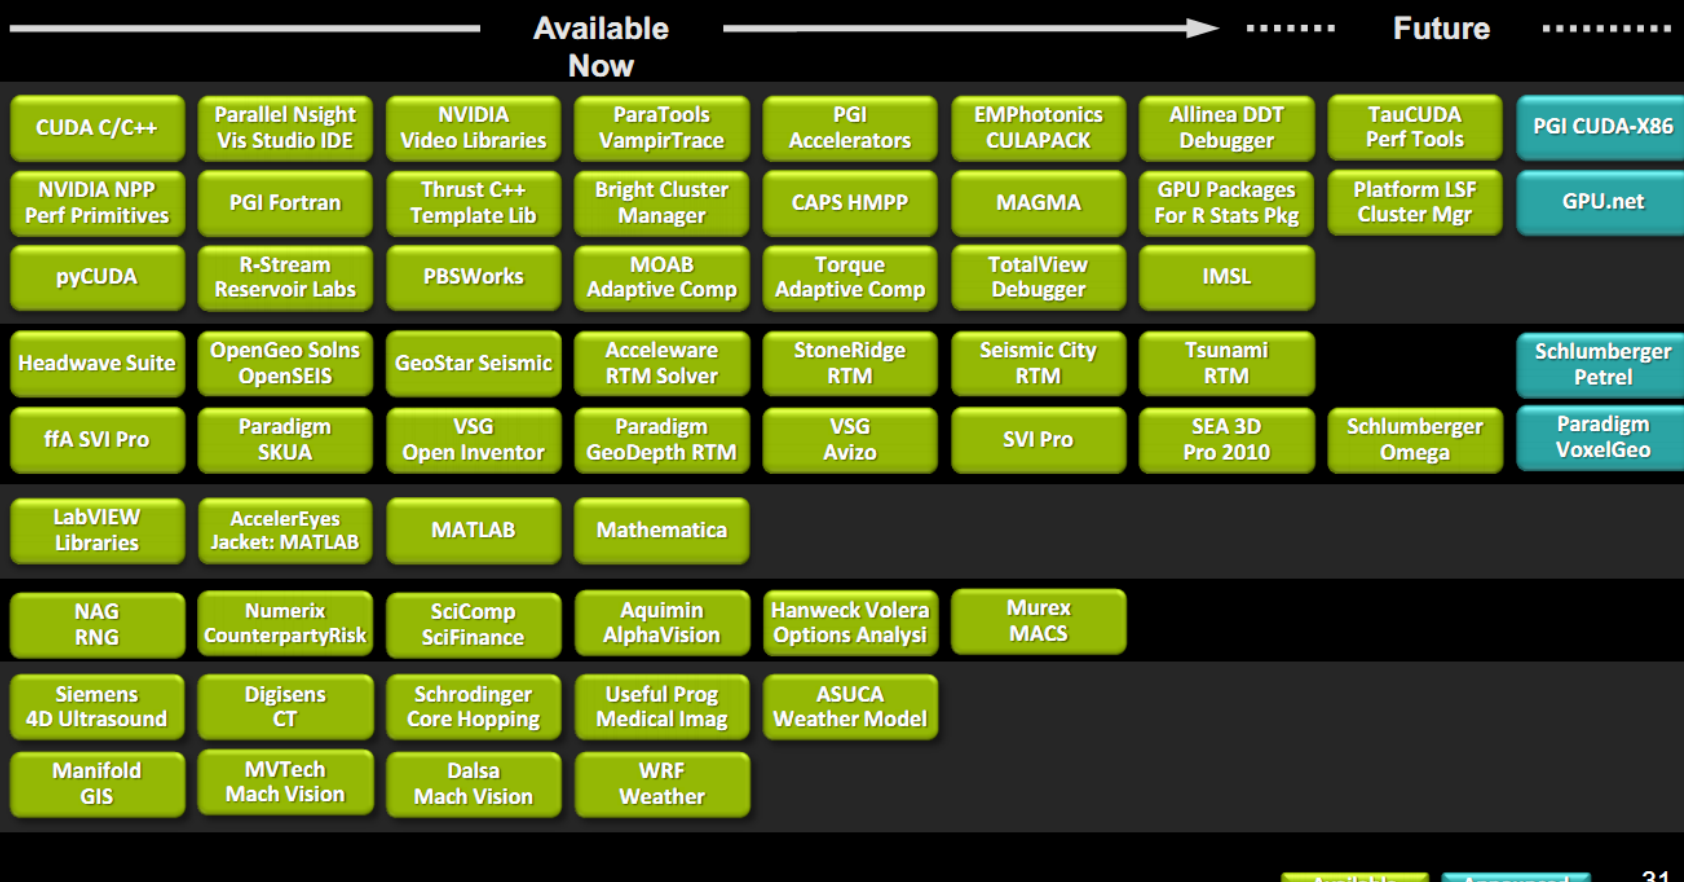

Increasing Number of Professional CUDA Applications

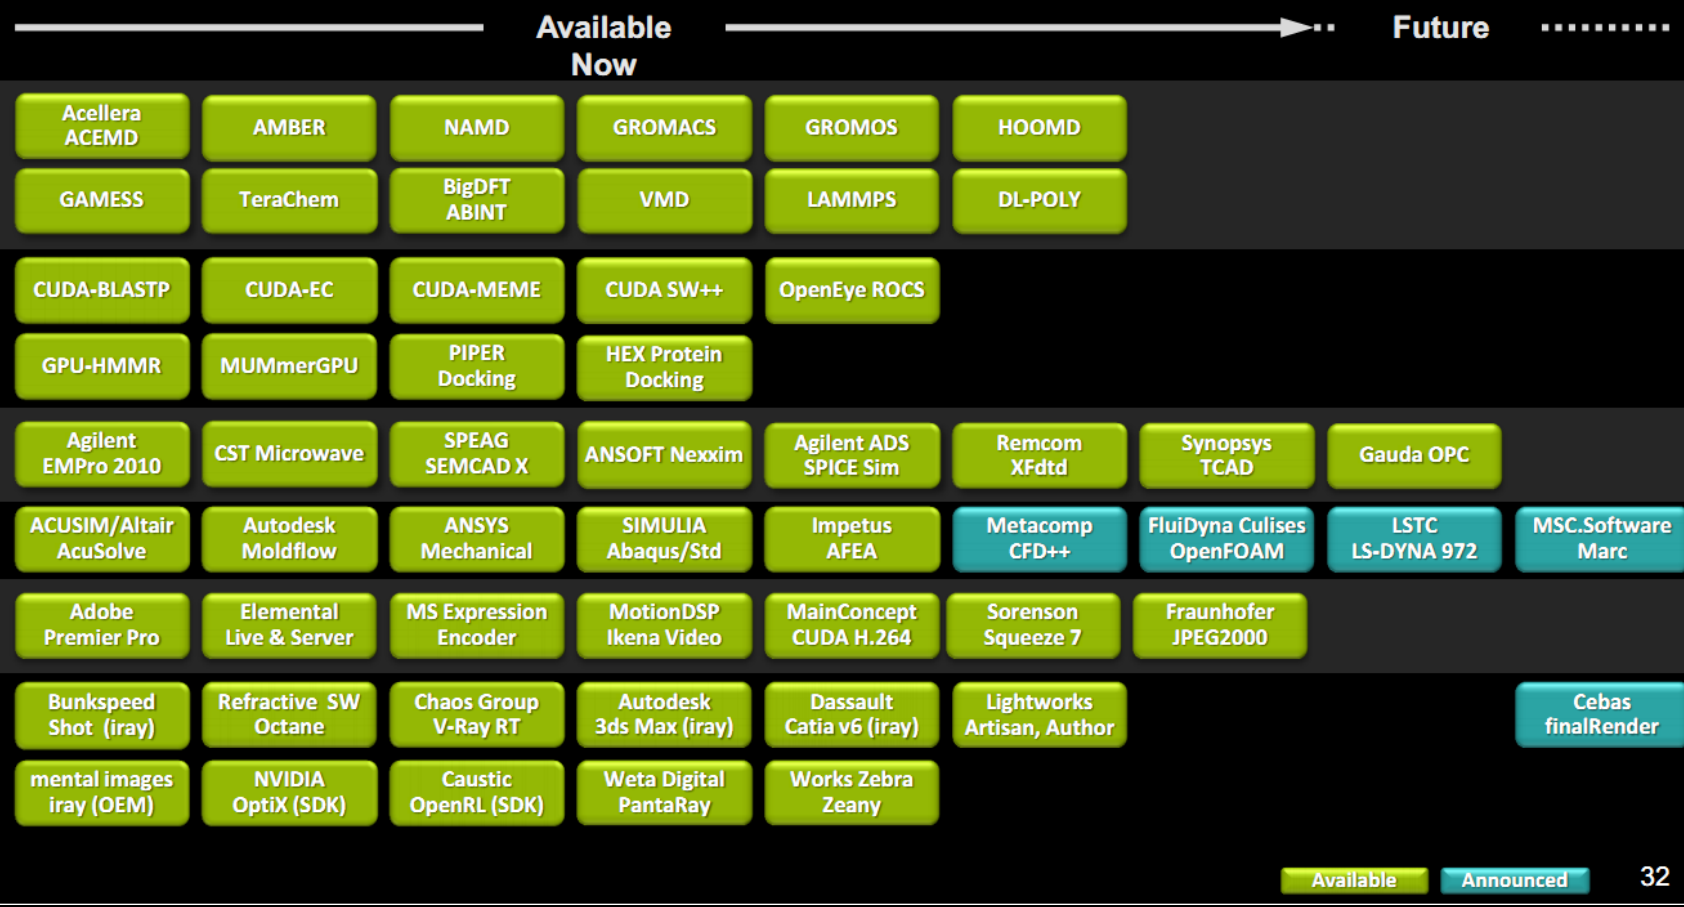

CUDA 4.0: Highlights

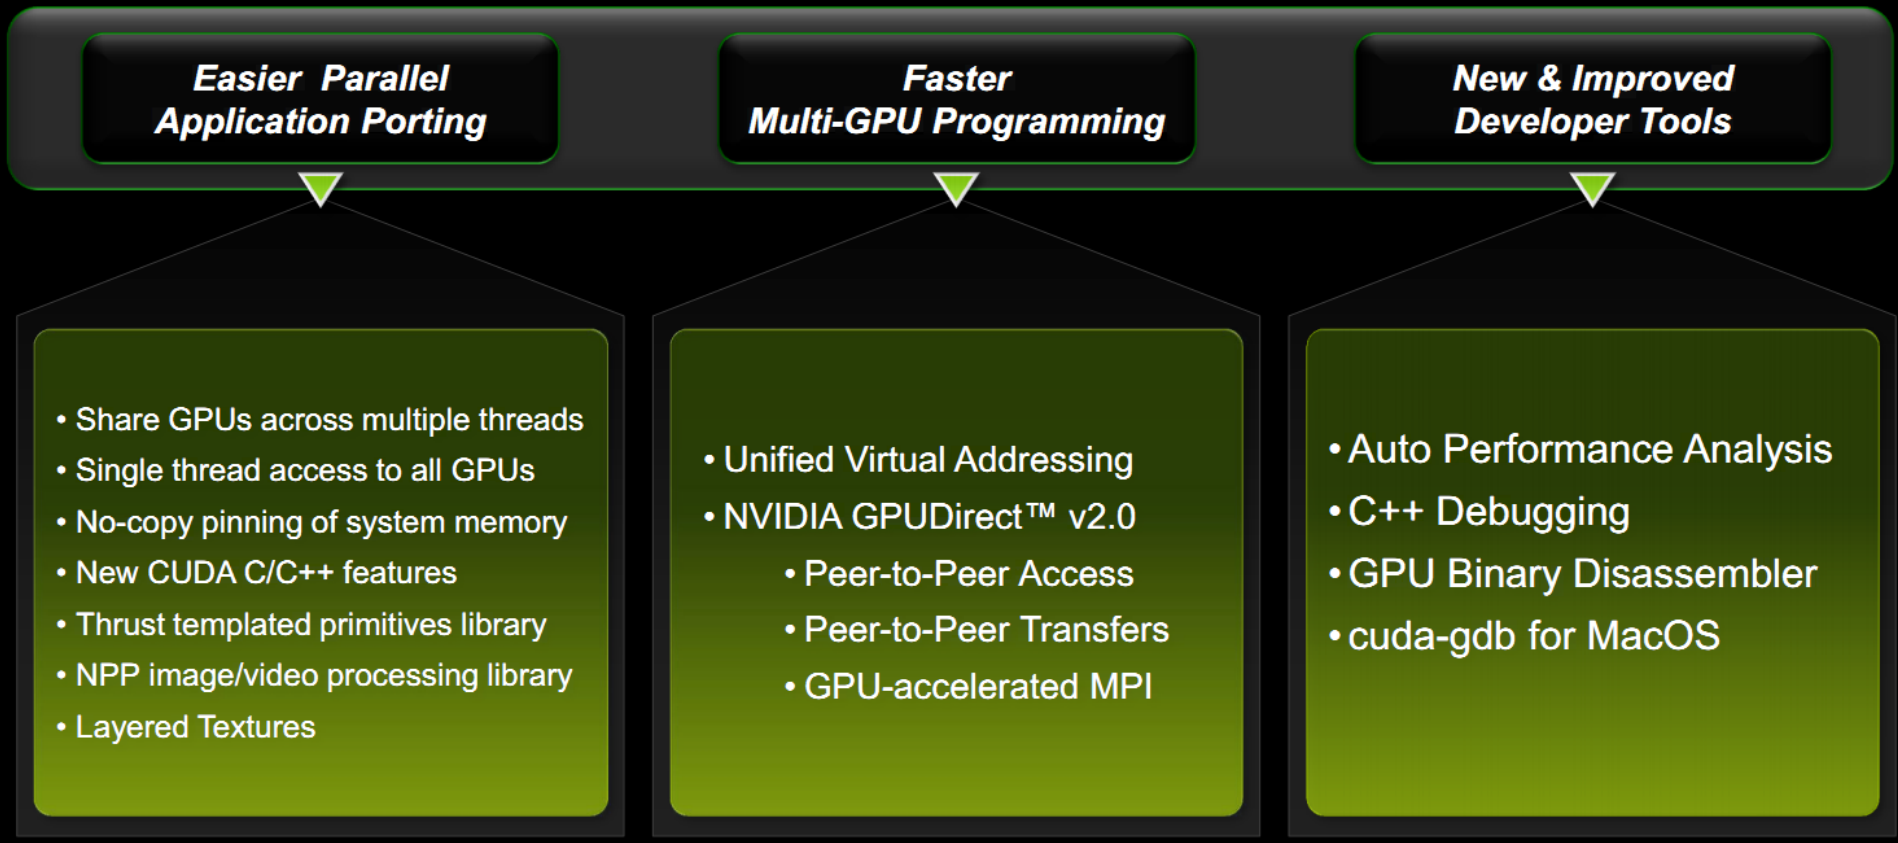

GPUDirect v2.0: Peer-to-Peer Communication

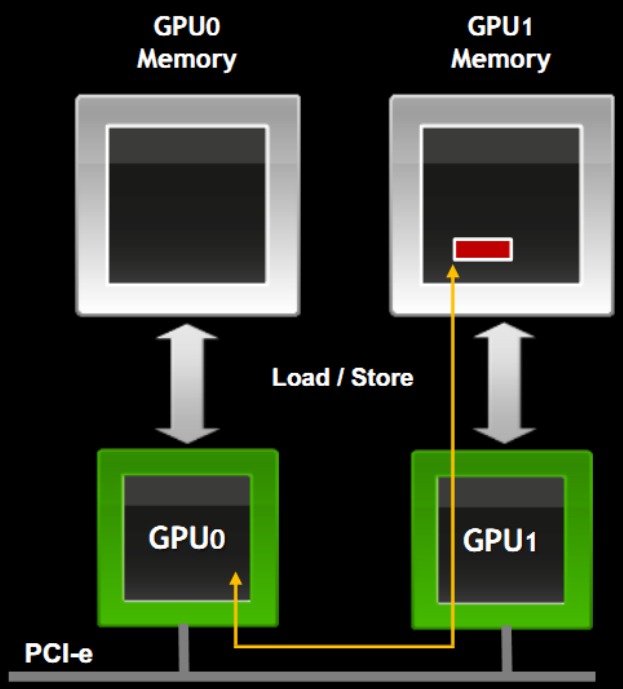

P2P Direct Access

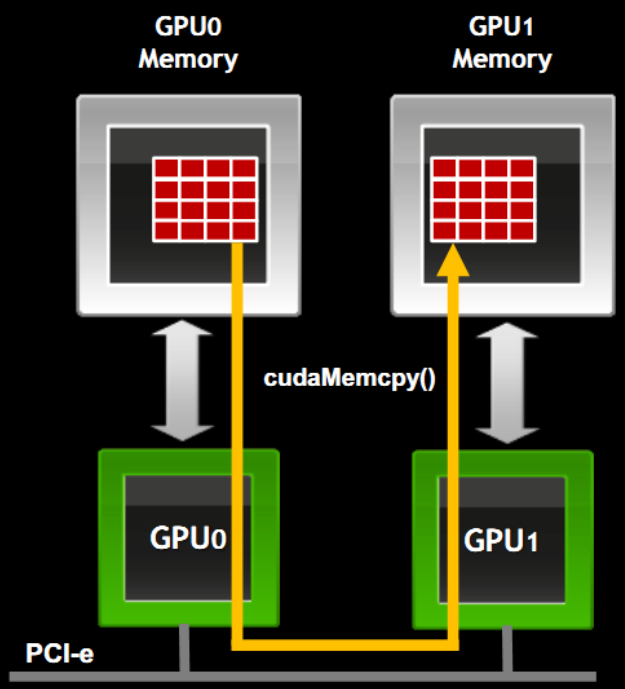

P2P Direct Transfers

Tesla Data Center & Workstation GPU Solutions

Tesla M-series GPUs
M2090 | M2070 | M2050

Servers & Blades

Tesla C-series GPUs
C2070 | C2050

Workstations

		M2090	M2070	M2050	C2070	C2050
Cores		512	448	448	448	448
Memory		6 GB	6 GB	3 GB	6 GB	3 GB
Memory bandwidth (ECC off)		177.6 GB/s	150 GB/s	148.8 GB/s	148.8 GB/s	148.8 GB/s
Peak Perf Gflops	Single Precision	1331	1030	1030	1030	1030
	Double Precision	665	515	515	515	515

GPU-Based OEM Systems

OEM	System	# of GPUs
Appro	Tetra	4GPUs in 1U
Asus	ESC4000	4GPUs in 2U
BULL	Bullx	18 GPUs in 7U
Dell	C61x	16 GPUs in 4U
HP	SL390G7	3 GPUs in 1U 8 GPUs in 2U
IBM	iDataplex	2 GPUs in 2U
NextIO	vCORE Express	4 GPUs in 1U
SGI	Prism XL	1 GPU in stick
Supermicro	6016GT-1U, TwinBlade-7126	2 GPUs in 1U 2 GPUs/blade
Tyan	FT72-B7015, GN70	8GPUs in 4U 3GPU/2U

Tesla M2090 vs M2070

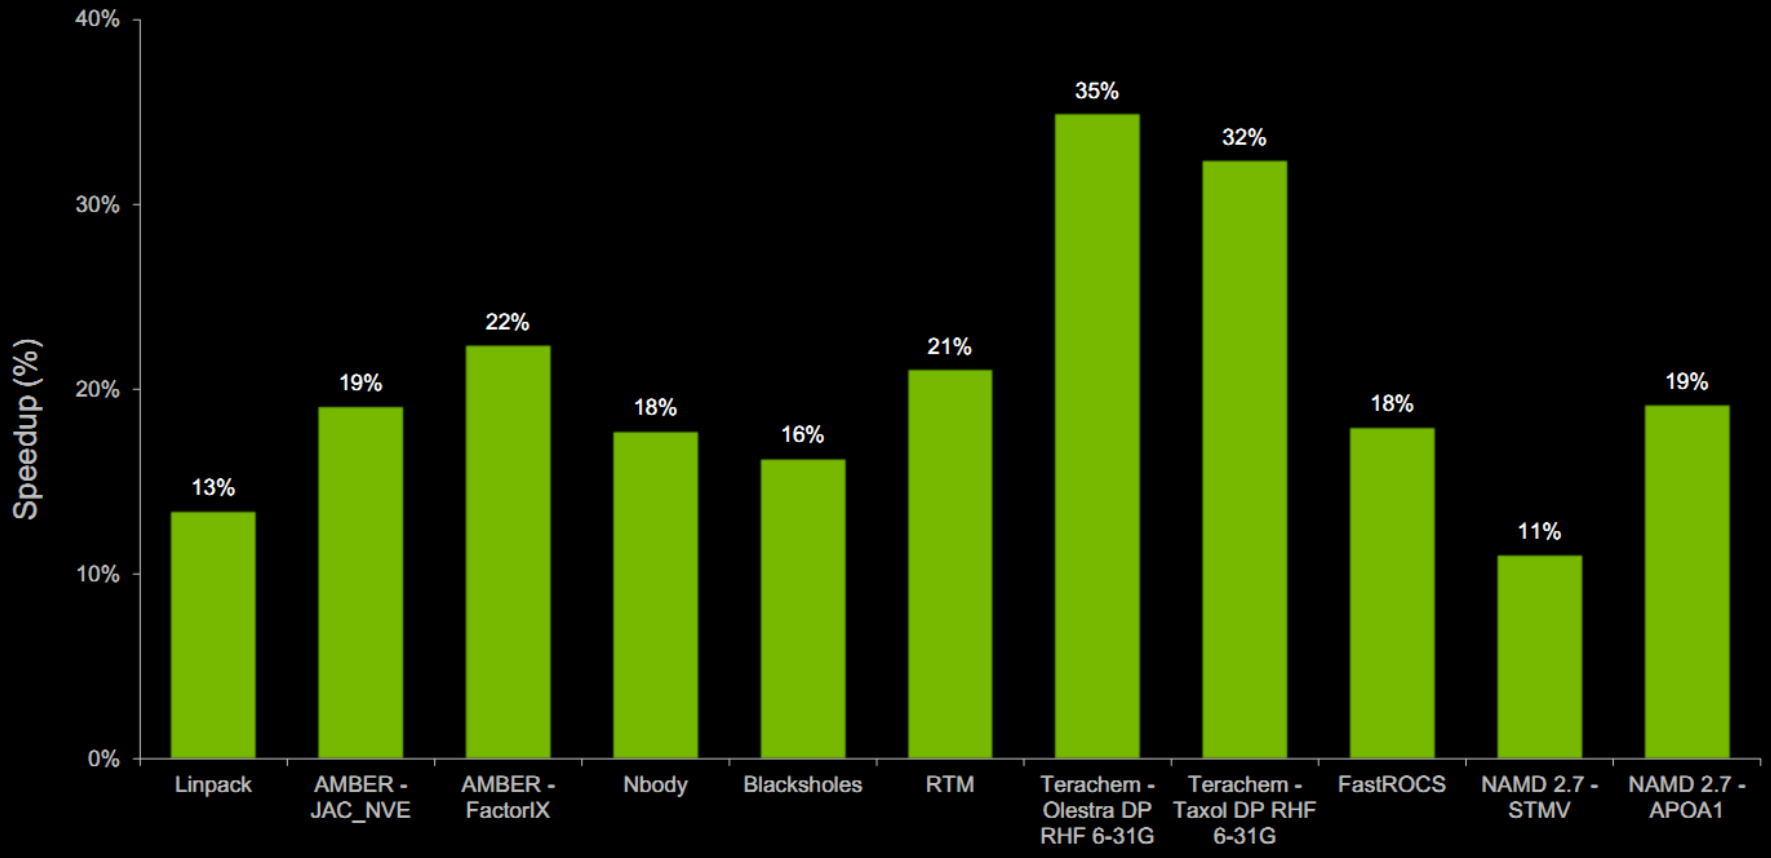

Maximize Compute Perf with Tesla M2090

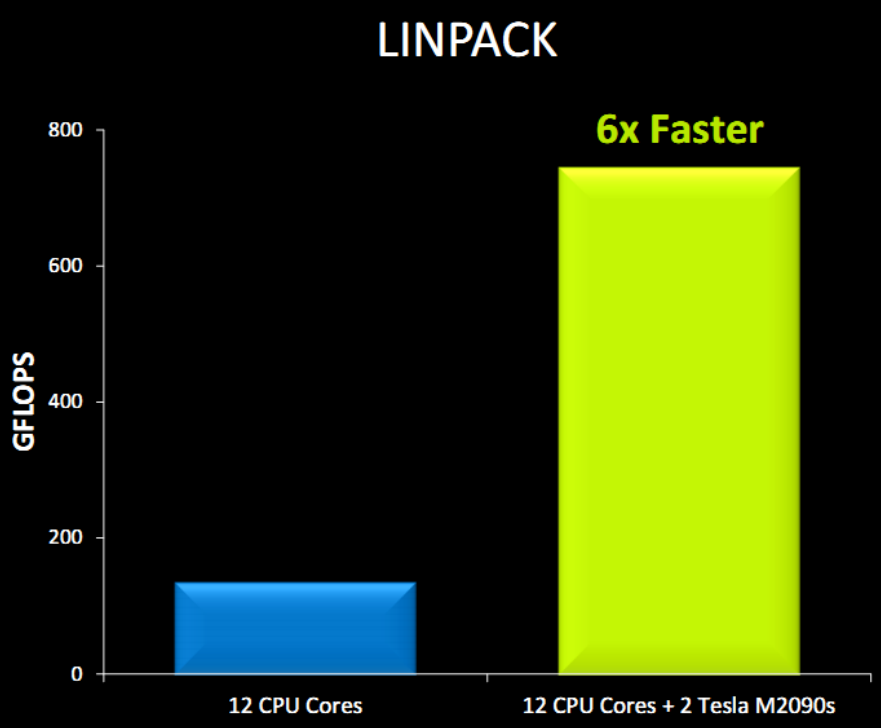

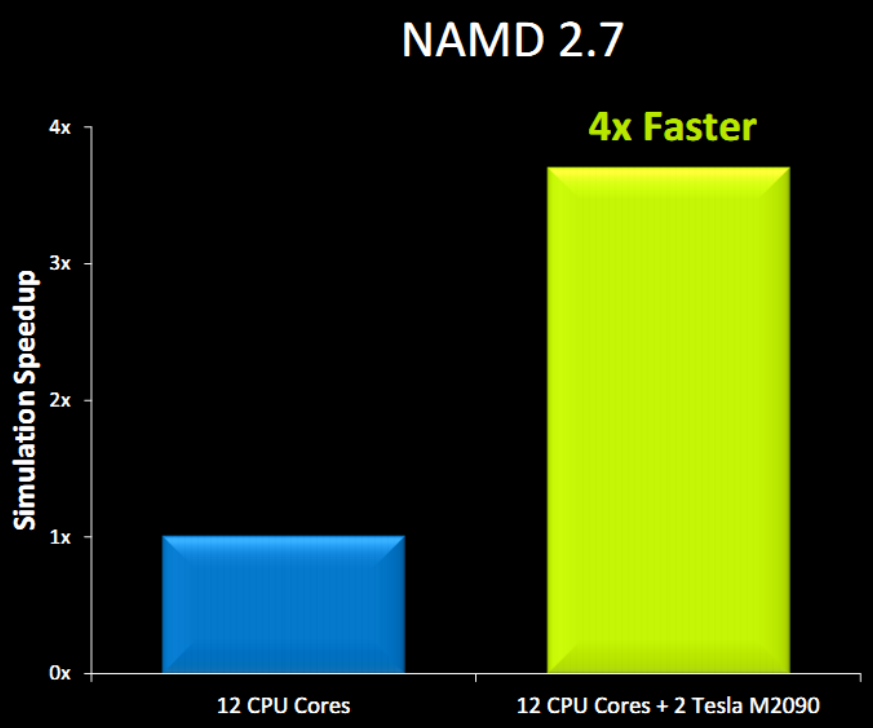

Benchmark
CPU: Dual socket Intel Xeon X5670, 2.93 GHz (12 cores)
Memory: 48 GB DDR3
CUDA 3.2

GPU: when & why?

4 CPU
0 GPU

High Density GPU

2 CPU
4 GPU

GPU optmized

4 CPU
4 GPU

Upgrade

4 CPU
4 GPU

GPU: cost/performance

x 3,26

x 3,27

High Density GPU

2 CPU
4 GPU

GPU optmized

4 CPU
4 GPU

x 2,91

x 2,83

4 CPU
0 GPU

Upgrade

x 3,65

x 2,81

4 CPU
4 GPU

Price

Price + PWR
+ cooling

Communication GPU to GPU

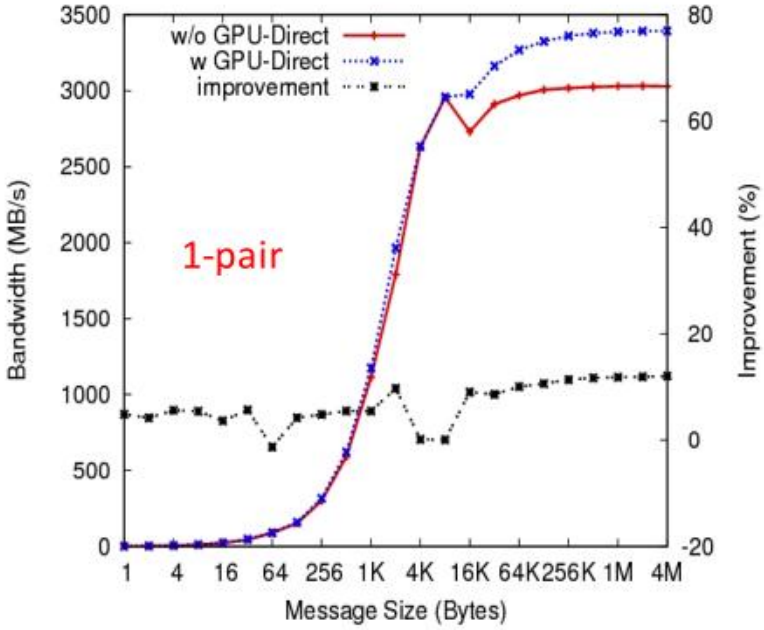

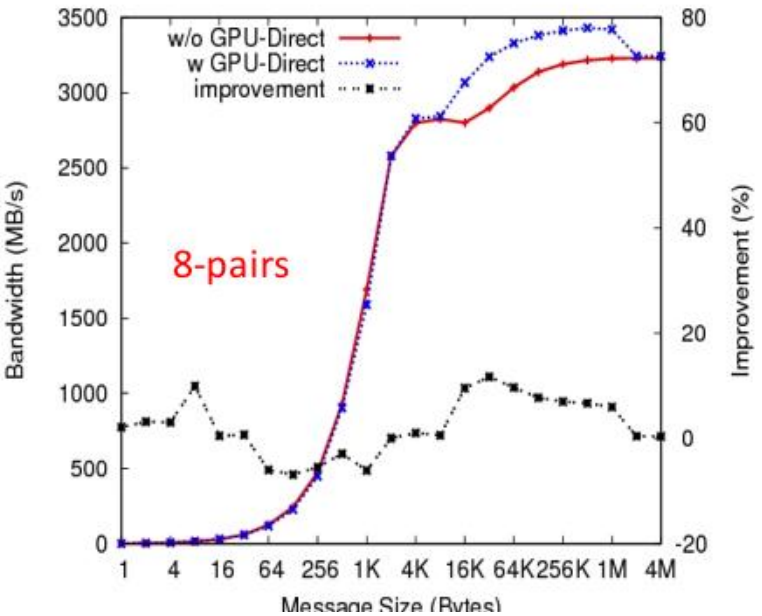

GPU direct: shared host buffer between GPU and HCA IB

Better performance up to 12%

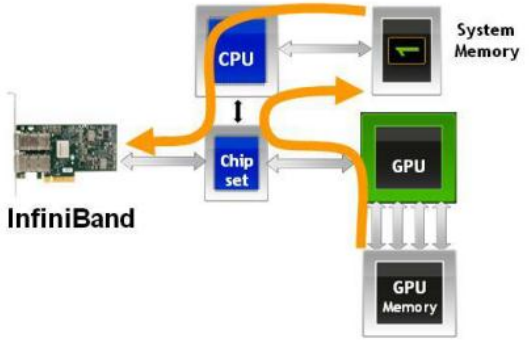

Fonte: Dhabaleswar K. (DK) Panda, Sayantan Sur (The Ohio State University) – Future of MPI, Tutorial at HPC Advisory Council Workshop, Lugano 2011

How to compare CPU and GPU calculations:

1. Never compare single core vs. single GPU

2. Take a look in the hardware configuration of the node

3. Comparison has to be done between the node with GPUs and without GPUs

Why NAMD?

Because you can define independently
the number of CPU cores & number of GPUs

```
>: Charmrun ++local +p 8 namd2 +idlepoll +devices 0,1 input.namd
```

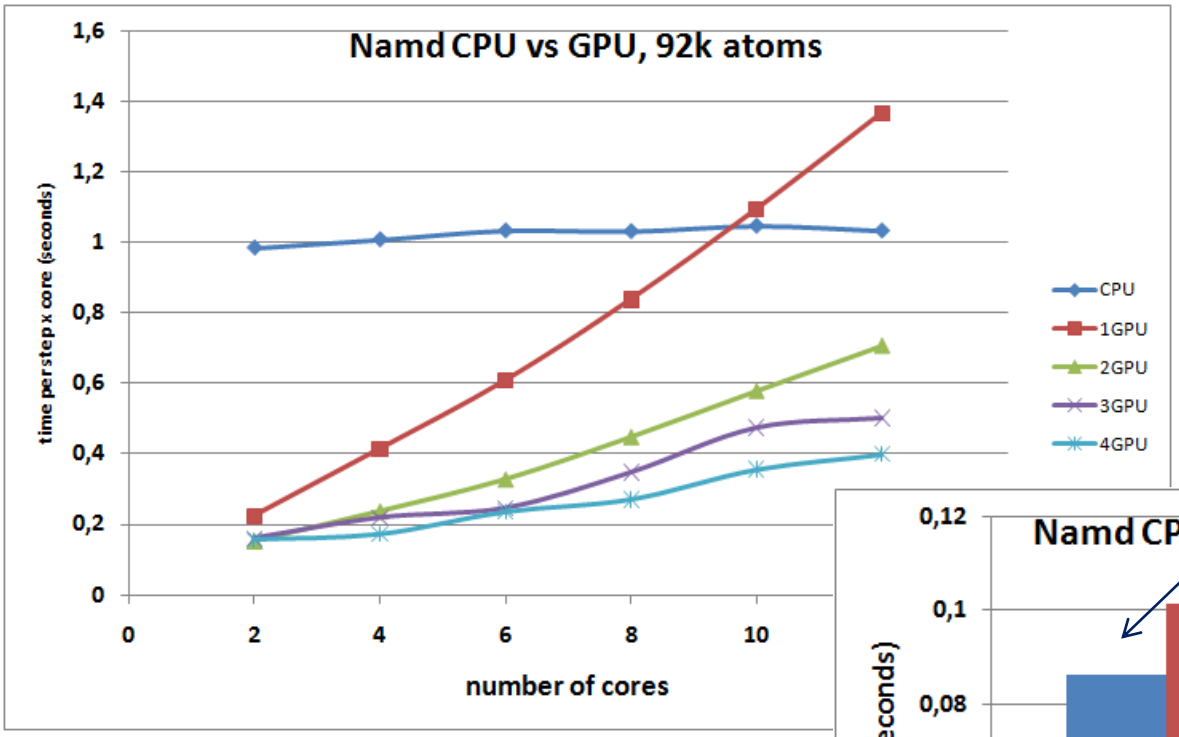
Namd CPU vs GPU, 92k atoms
time per step x core (seconds)
1,6
1,4
1,2
1
0,8
0,6
0,4
0,2
0
0
2
4
6
8
10
number of cores
CPU
1GPU
2GPU
3GPU
4GPU

Speedup 2.6x

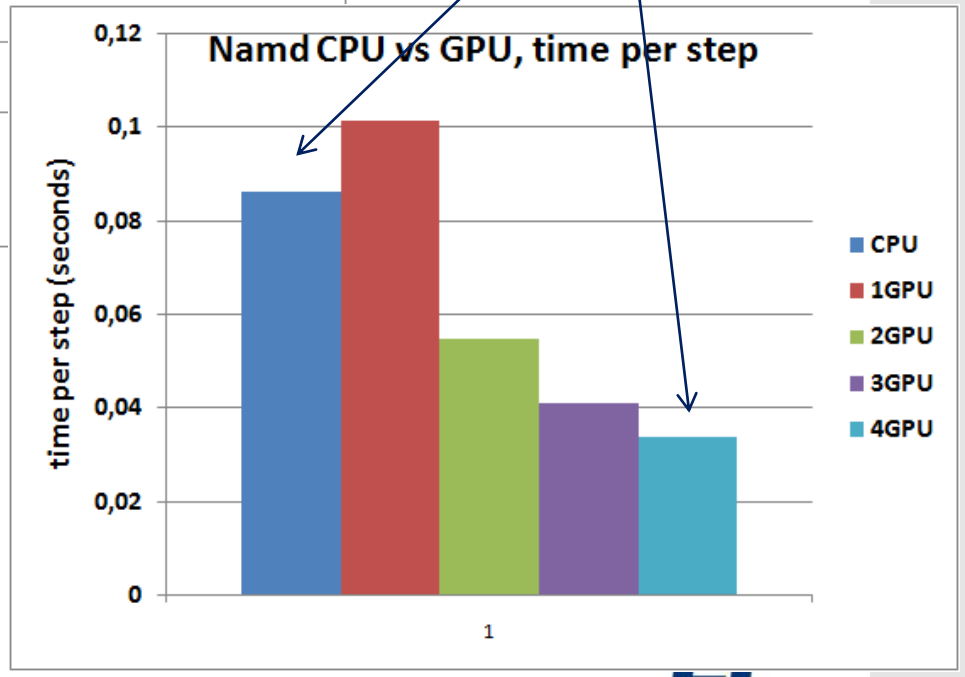
Namd CPU vs GPU, time per step
time per step (seconds)
0,12
0,1
0,08
0,06
0,04
0,02
0
1
CPU
1GPU
2GPU
3GPU
4GPU

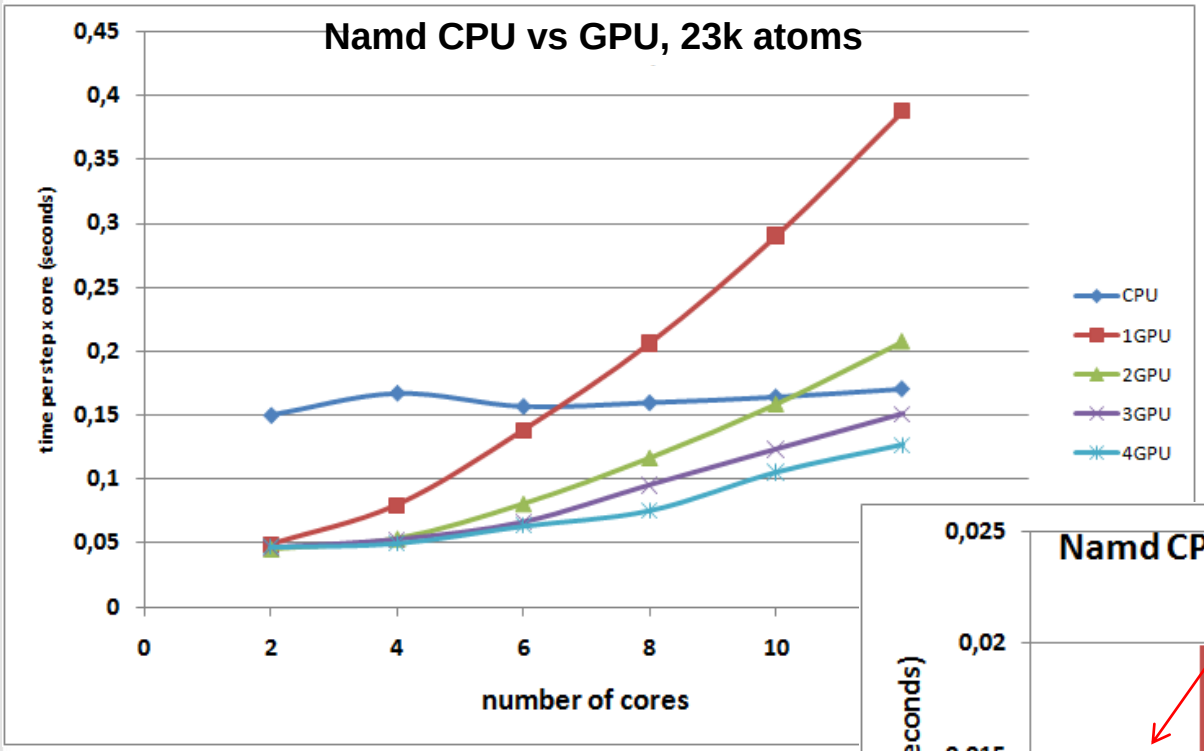
Namd CPU vs GPU, 23k atoms
time per step x core (seconds)
0,45
0,4
0,35
0,3
0,25
0,2
0,15
0,1
0,05
0
0
2
4
6
8
10
number of cores
CPU
1GPU
2GPU
3GPU
4GPU

NO SPEEDUP

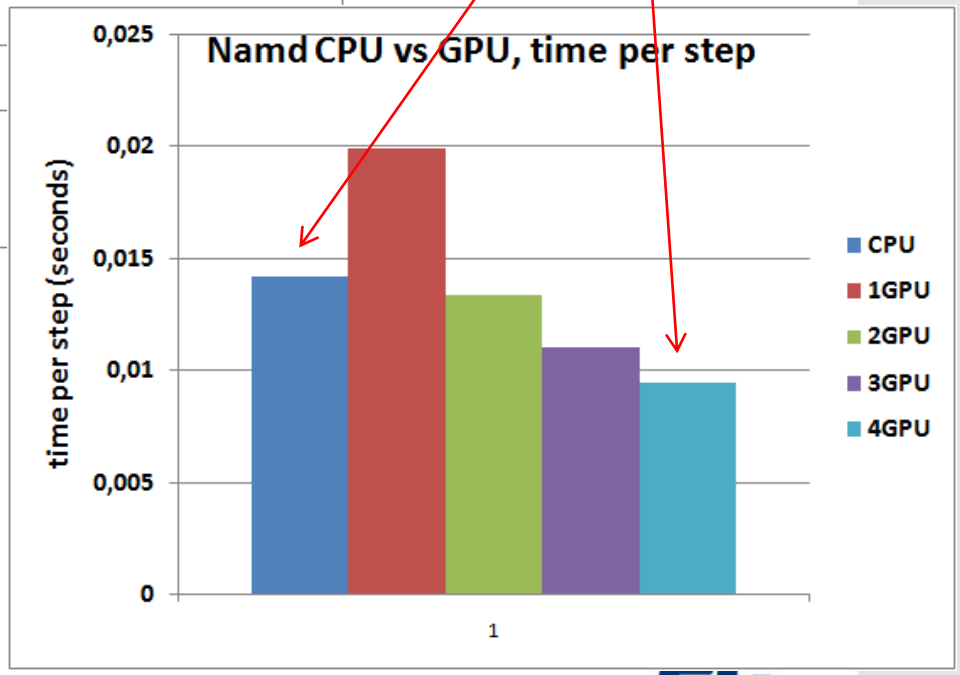
Namd CPU vs GPU, time per step
time per step (seconds)
0,025
0,02
0,015
0,01
0,005
0
1
CPU
1GPU
2GPU
3GPU
4GPU

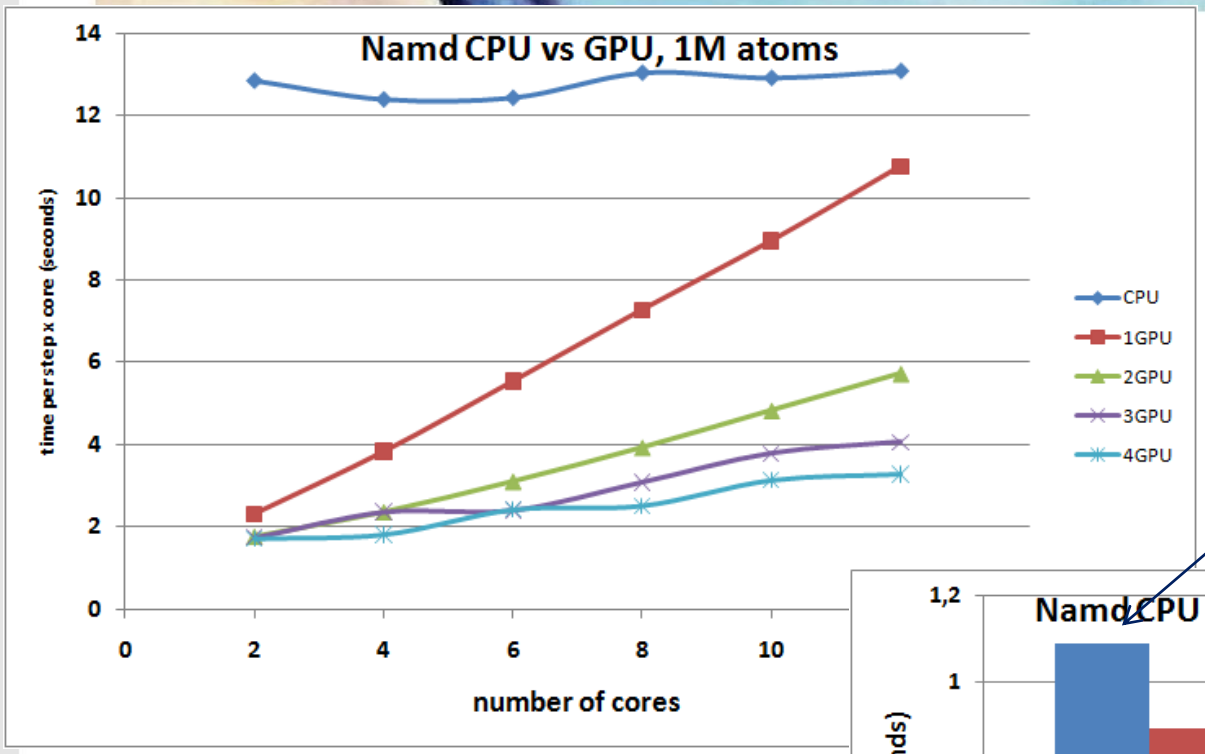
Namd CPU vs GPU, 1M atoms
time per step x core (seconds)
number of cores
CPU
1GPU
2GPU
3GPU
4GPU

Speedup 4x

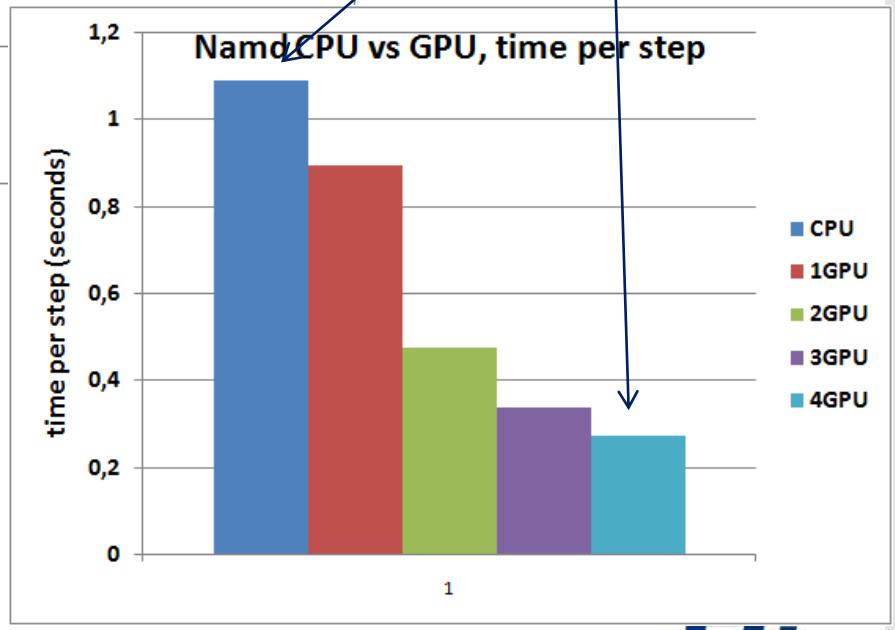
Namd CPU vs GPU, time per step
time per step (seconds)
CPU
1GPU
2GPU
3GPU
4GPU

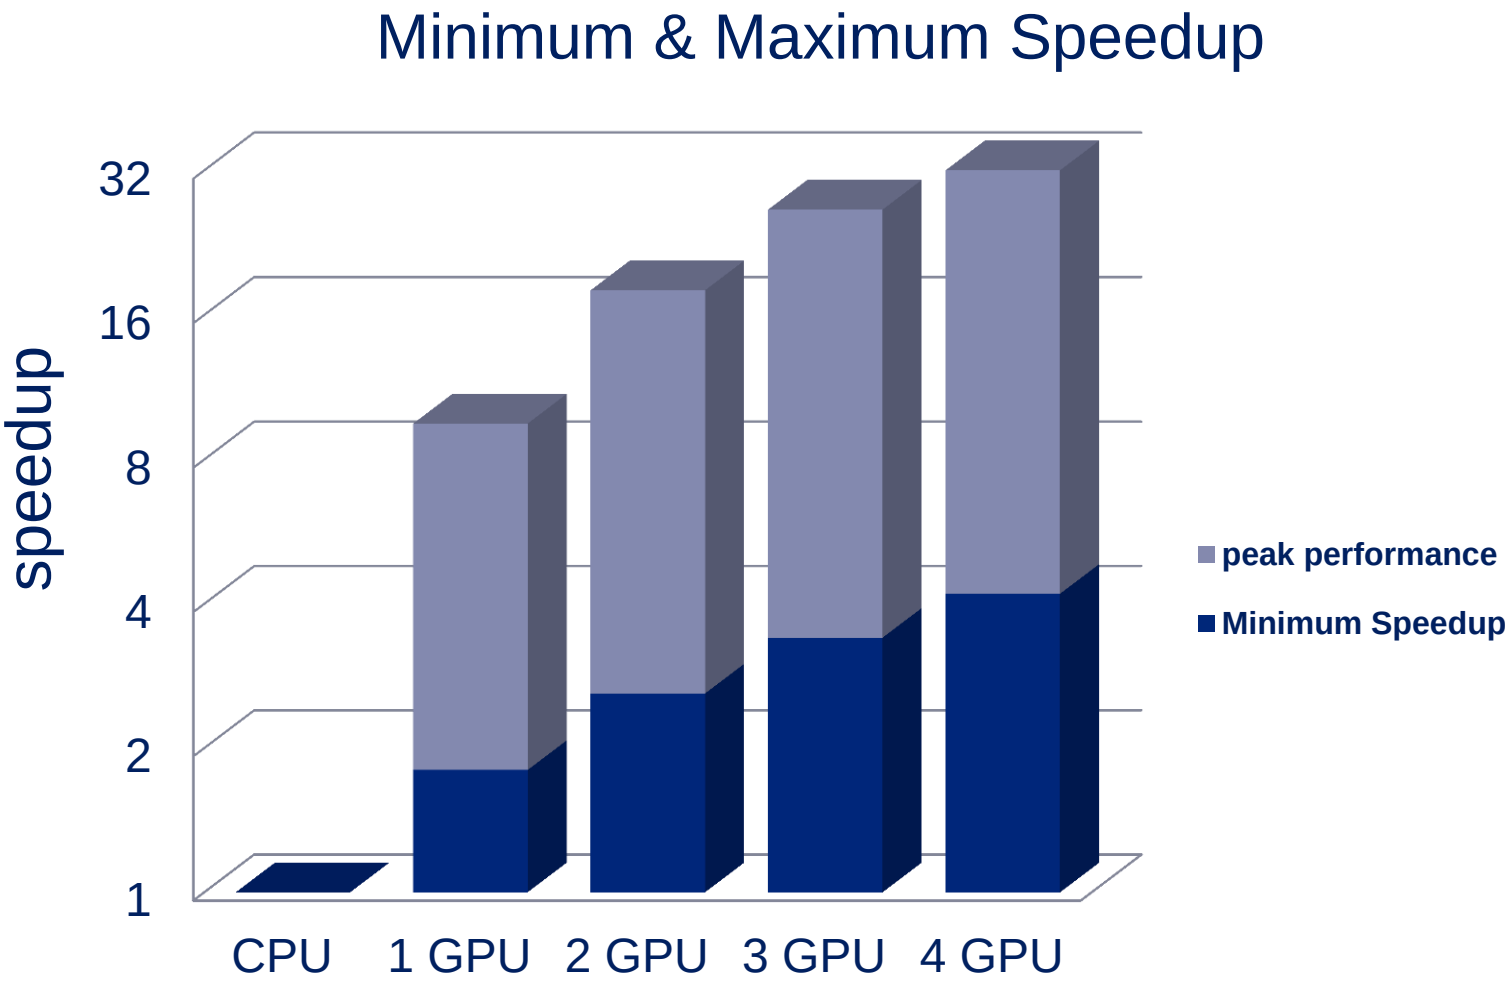
Minimum & Maximum Speedup
speedup
32
16
8
4
2
1
CPU
1 GPU
2 GPU
3 GPU
4 GPU
peak performance
Minimum Speedup

Features and Benefits

- Featuring up to four NVIDIA Tesla 20-Series GPUs per 1U server
- Supports up to two Quad/Six-Core Intel Xeon processor 5600 series per server
- Up to 96GB of DDR3 memory in 12 DIMM sockets
- Up to 3.0TB storage per server
- Two PCIe 2.0 x16 slots with riser card for GPU cards and one PCIe 2.0 x4 slot with riser
- Tool-less access to chassis, memory, HDDs, PCI card, blowers, and power supply
- Choice of Linux or Windows operating systems
- Offers a flexible, reliable and scalable compute platform
- Easy to maintain and service – hot-swappable drives and fans
- Energy-efficient - 1400W high-efficiency power supply and twelve cooling fans
- Best TCO with improved system density, computing capability while keeping the datacenter cool

server E7228+Tesla™

The new supercomputing technology

TECHNICAL SPECIFICS TESLA™ S2050

2 x Tesla S2050 each one of which with the following features:

- 4 x GPU Tesla™
- 4 Teraflops in Single Precision
- 2 Teraflops in Double Precision
- GPU processor Clock 1.55 GHz peak clock
- Memory: up to 24 GB DDR5
- System I/O: two PCIe connection
  Each connection leads to two of 4 GPUs

TECHNICAL SPECIFICS E7228

4 Motherboards Dual Socket Intel® Xeon® Six-Core serie 5600, each one of them with the following features:

- 1 x chipset Intel® 5520, 1333/1066/800 MHz
- 1 x 96GB REG ECC DDR3 1333/1066/800MHz (12 DIMMs)
- 1 x PCI-Express 16x 2.0 (Low Profile)
- 3 x SAS/SATA drive bays (2 bays)
- 1 x Intel® ICH10R – SATA II SW Raid - 6 ports (integrated)
- 1 x LAN Intel® Dual Gigabit 82576 controllers
  (2 LAN ports total)
- 1 x Matrox G200eW
- 1 x sets of rear I/O ports including 2 USB 2.0, VGA, COM
- 1 x sets of IPMI 2.0 with dedicated Realtek 10/100Mb/s LAN port, KVM-over-LAN

E7228 is also equipped with:

- 2 x 1400W PFC power supply
  Independent power control each node has its own power management

workstation E7095

PERSONAL SUPERCOMPUTER
NVIDIA® TESLA™

E7095 TECHNICAL FEATURES

- Workstation tower/rackmount low noise with optimized cooling and hetsink
- 8 bays SAS/SATA hot swap
- 1400W redundant power supply
- M/B dual Intel® Xeon®
- Chipset Intel® 5500
- CPU Intel® Xeon® serie 5600
- Up to 96 GB memory DDR3 reg ECC
- DVD writer dual layer
- Up to 4 x Tesla C1060/C2050 or 2x S1070/S2050
- Graphic adapter NVIDIA Quadro NVS o FX

Thank you

E4 Computer Engineering SpA

Via Martiri della Liberta’ 66

42019 - Scandiano (RE), Italy

www.e4company.com

Switchboard: +39.0522.991811

QUESTIONS?

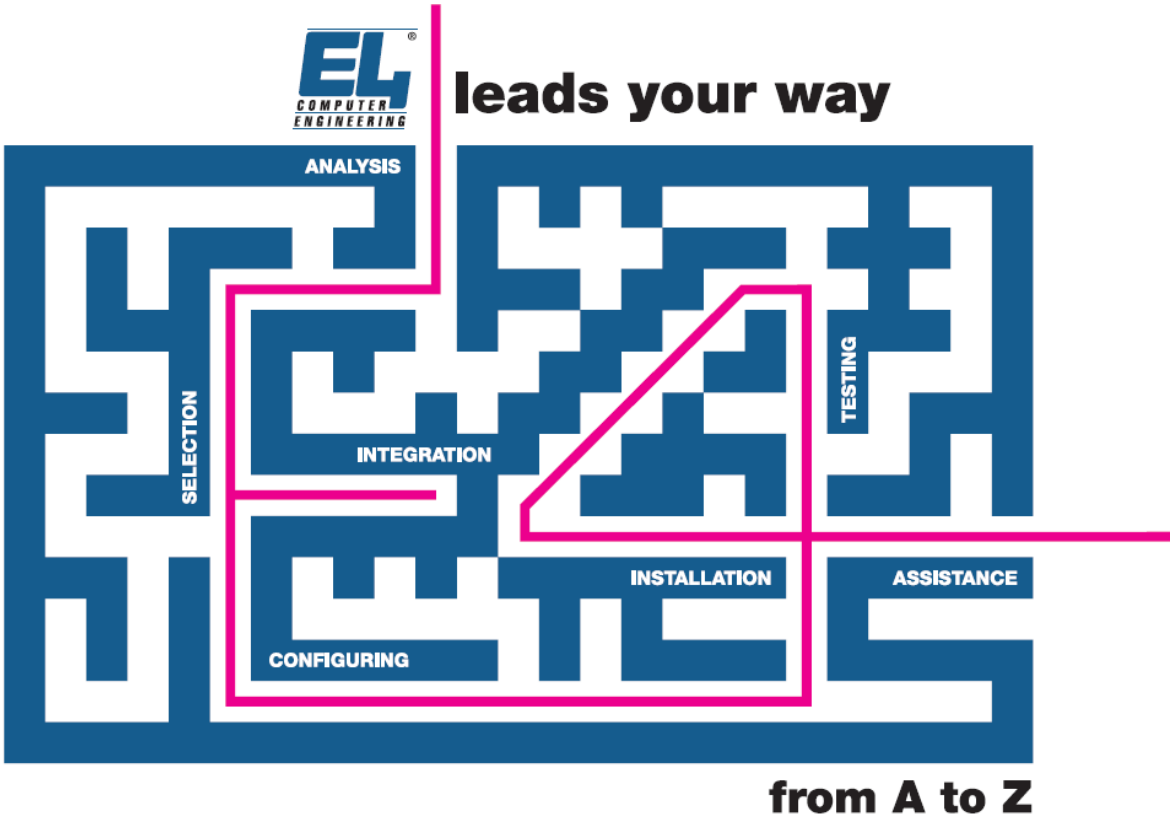
EL
COMPUTER ENGINEERING
leads your way
ANALYSIS
SELECTION
INTEGRATION
TESTING
INSTALLATION
ASSISTANCE
CONFIGURING
from A to Z