Kepler GPU Architecture and Benefits to Earth System Modeling

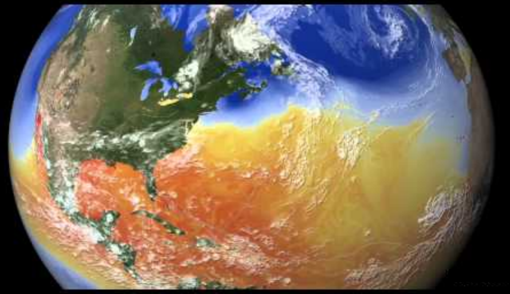

Stan Posey | HPC Industry and Applications Development| sposey@nvidia.com

Agenda: Kepler GPU Architecture and Benefits to Earth System Modeling

- Introduction of GPUs in HPC
- Review of GPU Progress
- Kepler Architecture Benefits

GPU Computing is Mainstream HPC

NVIDIA

Oil & Gas

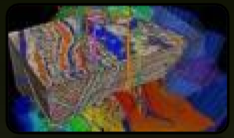

- Schlumberger
- BR Petrobras
- Total
- Chevron
- Paradigm

Edu/Research

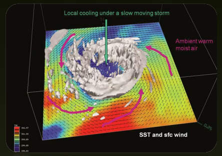

- Chinese Academy of Sciences
- Georgia Tech
- Harvard School of Engineering and Applied Sciences
- Oak Ridge National Laboratory

Government

- Air Force Research Laboratory
- NASA
- Naval Research Laboratory
- BAE Systems

Life Sciences

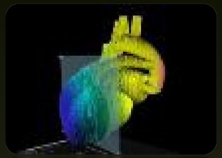

- Boston Scientific
- MGH 1811 Mass General Hospital
- Max Planck Institute
- Beckman Coulter

Finance

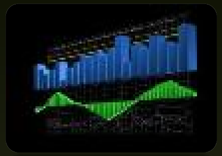

- Bloomberg
- BNP Paribas
- Standard Life
- J.P.Morgan
- NumeriX

Manufacturing

- Altair
- ANSYS
- Autodesk
- 3DS Simulia
- MSC Software

OEM Servers with NVIDIA GPUs Available Today

NVIDIA

Bull

hp

IBM

sgi

CRAY THE SUPERCOMPUTER COMPANY

T-PLATFORMS

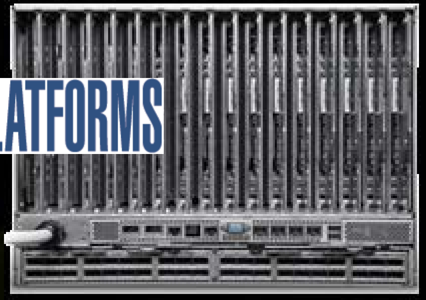

NextIO

SUPERMICRO

APPRO

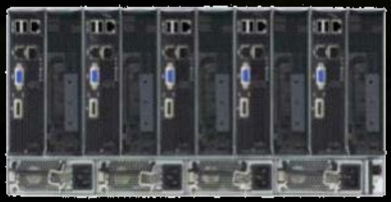

DELL

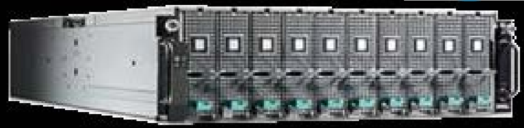

TYAN

Supercomputing Momentum With GPUs

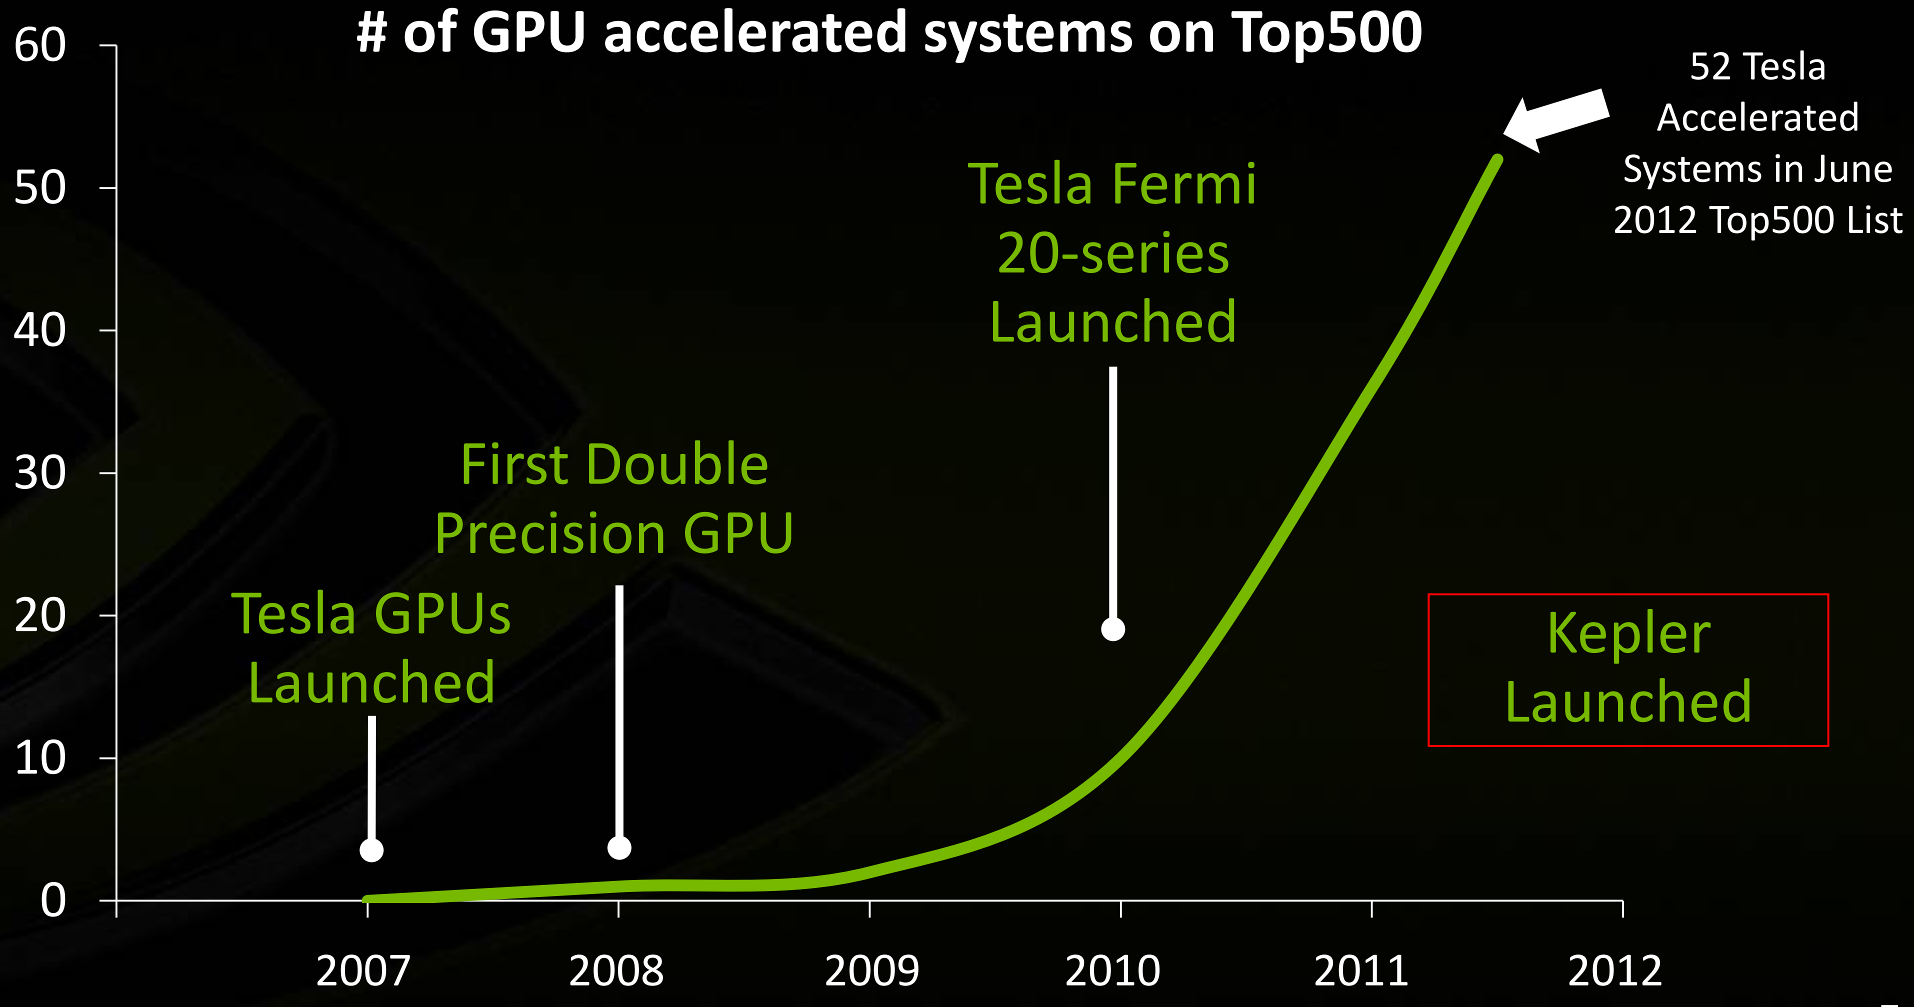

GPU Motivation: Performance and Power Efficiency

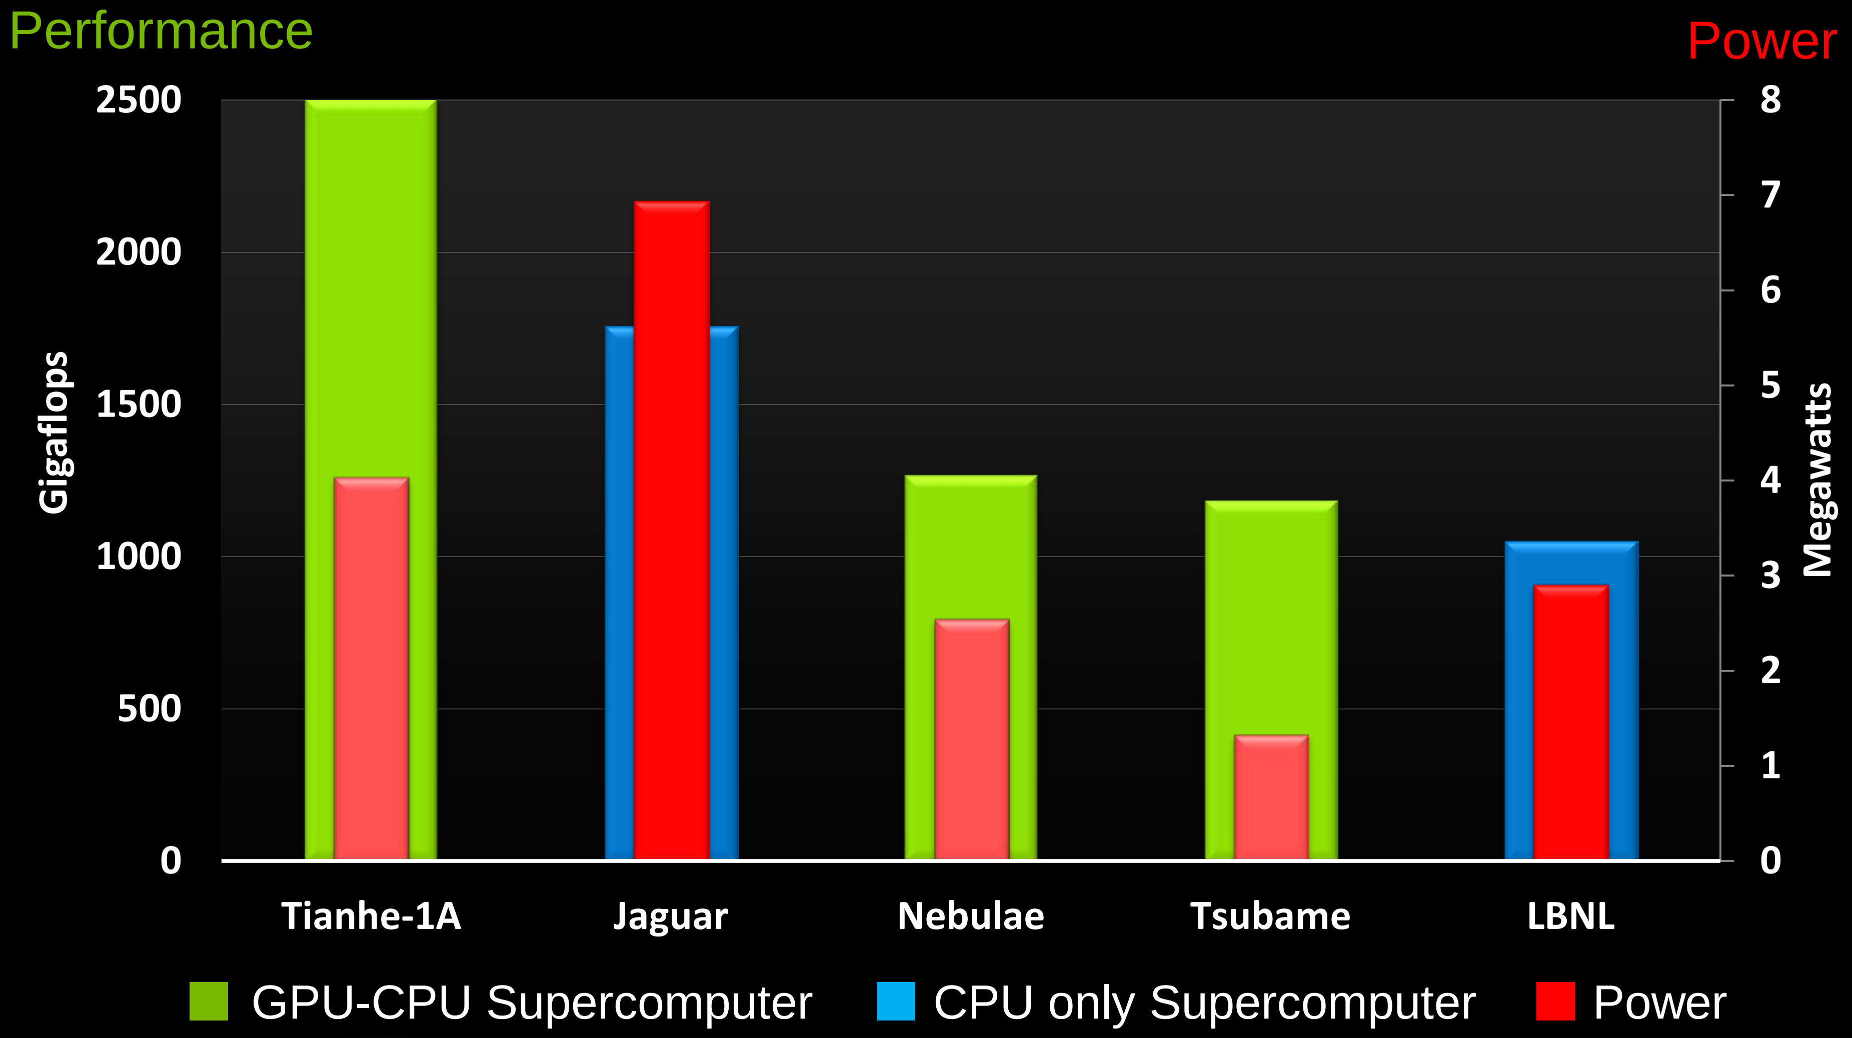

TiTech Winner of 2011 Gordon Bell Prize

Achieved with NVIDIA Tesla GPUs

Special Achievement in Scalability and Time-to Solution

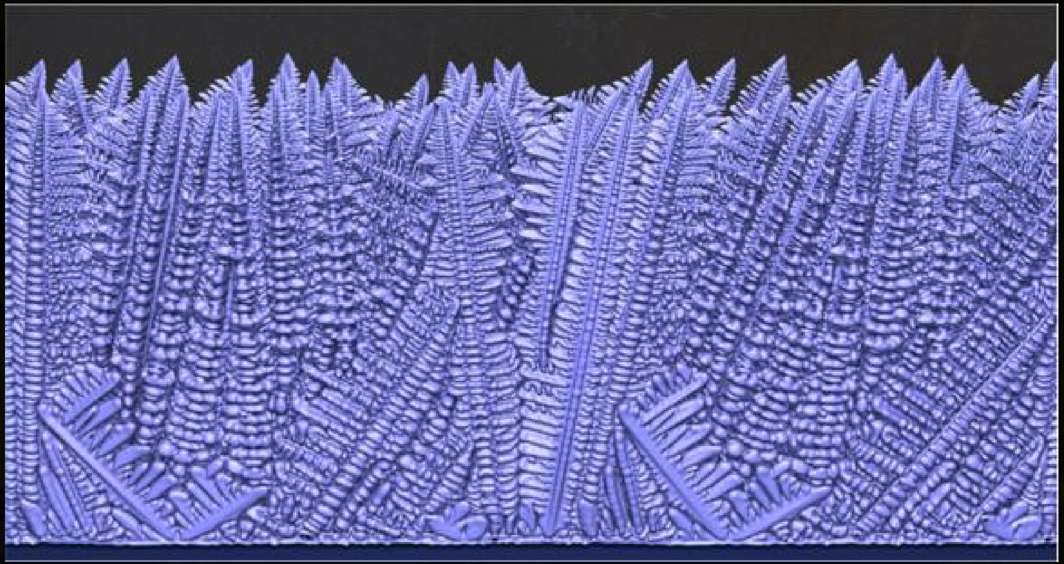

"Peta-scale Phase-Field Simulation for Dendritic Solidification on the TSUBAME 2.0 Supercomputer"
-- T. Shimokawabe, T. Aoki, et. al.

Tsubame 2.0
Tokyo Institute of Technology

4,224 Tesla GPUs + 2,816 x86 CPUs

ASUCA and NWP Achievement: 145 TFLOPS

Tsubame 2.0

東京工業大学
Tokyo Institute of Technology

Tokyo Institute of Technology

- 1.19 Petaflops
- 4,224 Tesla M2050 GPUs

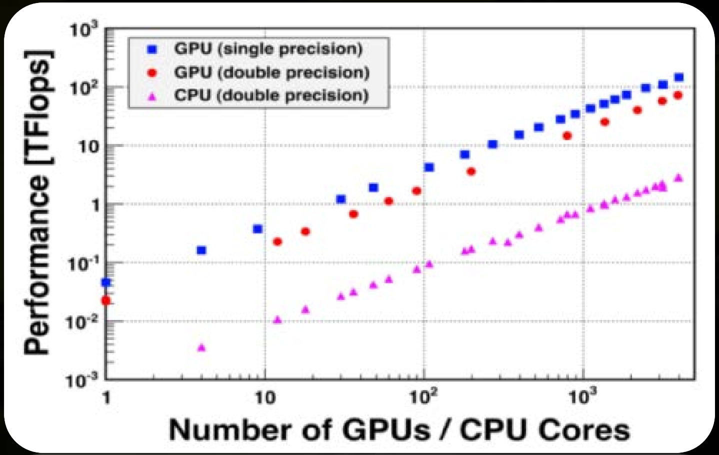

3990 Tesla M2050s

145.0 Tflops SP

76.1 Tflops DP

ASUCA and NWP Simulation on Tsubame 2.0, TiTech Supercomputer:
Dr. Takayuki Aoki, GSIC, Tokyo Institute of Technology, Tokyo Japan

Agenda: Kepler GPU Architecture and Benefits to Earth System Modeling

- Introduction of GPUs in HPC
- Review of GPU Progress
- Kepler Architecture Benefits

NVIDIA Technology and Strategy

Technology

- **Development of GPUs as a co-processing accelerator for x86 CPUs**
  - GPUs in-line with trend of higher resolution modeling with manageable compute costs

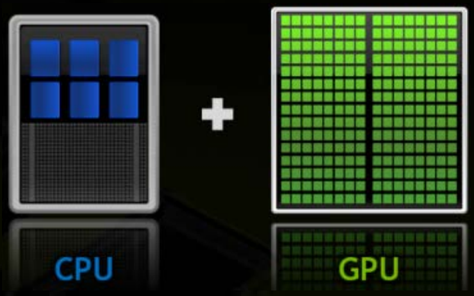

Strategy

- **Alliances to develop Fortran-based GPU compilers and tools**
  - Commercial vendors PGI, CAPS, and Cray; OpenACC membership; Research initiatives
- **Customer collaborations in applications engineering**
  - NVIDIA technical contributions to CAM-SE, COSMO, WRF, and NEMO
- **GPU integration with systems from all major vendors**
  - IBM, Cray, HP, SGI and many others; Kepler based-systems available during 4Q 2012
  - Technical collaboration on large system projects such Titan/ORNL, TSUBAME/Titech, etc.

Example: NASA and Global Cloud Resolving GEOS-6

Programming weather, climate, and earth-system models on heterogeneous multi-core platforms

September 7-8, 2011 at the National Center for Atmospheric Research in Boulder, Colorado

The Finite-Volume Dynamical Core on GPUs within GEOS-5

- Dr. William Putman, Global Modeling and Assimilation Office, NASA GSFC

NASA targeting GEOS global model resolution at sub-10-km to 1-km range

Computational requirements for typical 5-day operational forecast:

Grid resolution	Westmere CPU cores	Comments
10 KM	12,000	Possible today
3 KM	300,000	Reasonable but not available
1 KM	10,000,000	Impractical, need accerators

3.5-km GEOS-5 Simulated Clouds (CPU-Only)

Source: http://data1.gfdl.noaa.gov/multi-core/

GPU Progress Reported at NCAR Workshop

NVIDIA

Programming weather, climate, and earth-system models on heterogeneous multi-core platforms

September 7-8, 2011 at the National Center for Atmospheric Research in Boulder, Colorado

GPU related talks (11+) that cover application software such as:

NIM | WRF | GEOS-5 | HOMME | COSMO | CAM-SE | ICON

Successes and Challenges using GPUs for Weather and Climate Models	Mark Govett, NOAA
Experience using FORTRAN GPU Compilers with the NIM	Tom Henderson, NOAA
GPU Acceleration of the RRTM in WRF using CUDA FORTRAN	Greg Ruetsch, NVIDIA
Lessons Learned adapting GEOS-5 GCM Physics to CUDA FORTRAN	Matt Thompson, NASA
Accelerated Cloud Resolving Model in Hybrid CPU-GPU Clusters	Jose Garcia, NCAR
Reworking Boundary Exchanges in HOMME for Many-Core Nodes	Ilene Carpenter, NREL
Performance optimizations for running an NWP model on GPUs	Jacques Middlecoff, NOAA
Rewrite of the COSMO Dynamical Core	Mueller / Gysi, SCS/CSCS
Experiences with the Finite-Volume Dynamical core and GEOS-5 on GPUs	Bill Putman, NASA
Progress in Accelerating CAM-SE	Jeff Larkin, Cray/ORNL
Porting the ICON Non-hydrostatic Dynamical Solver to GPUs	Will Sawyer, CSCS

Source: http://data1.gfdl.noaa.gov/multi-core/

GPU Considerations for ES Modeling

- Initial efforts are mostly dynamical core developments
  - If dynamics ~50% of profile time – 2x overall speed-up is possible
  - *Implicit schemes* – iterative sparse matrix linear algebra solvers
  - *Explicit schemes* – no linear algebra, operations on i,j,k stencil

- Increasing use of GPU-based libraries and directives
  - Examples: SpMV for implicit; stencil libraries, OpenACC directives

- Most models use a domain decomposition parallel method
  - Fits GPU model very well and preserves costly MPI investment

- Fortran programming on GPUs most critical for adoption
  - NVIDIA investments in CAPS, PGI and Cray compilers; OpenACC

Current Implementations Focus on Dynamics

GPU Strategies:

- Hand-CUDA (CUDA C, CUDA Fortran)
- GPU Libs (Stencils, etc.)
- OpenACC Directives

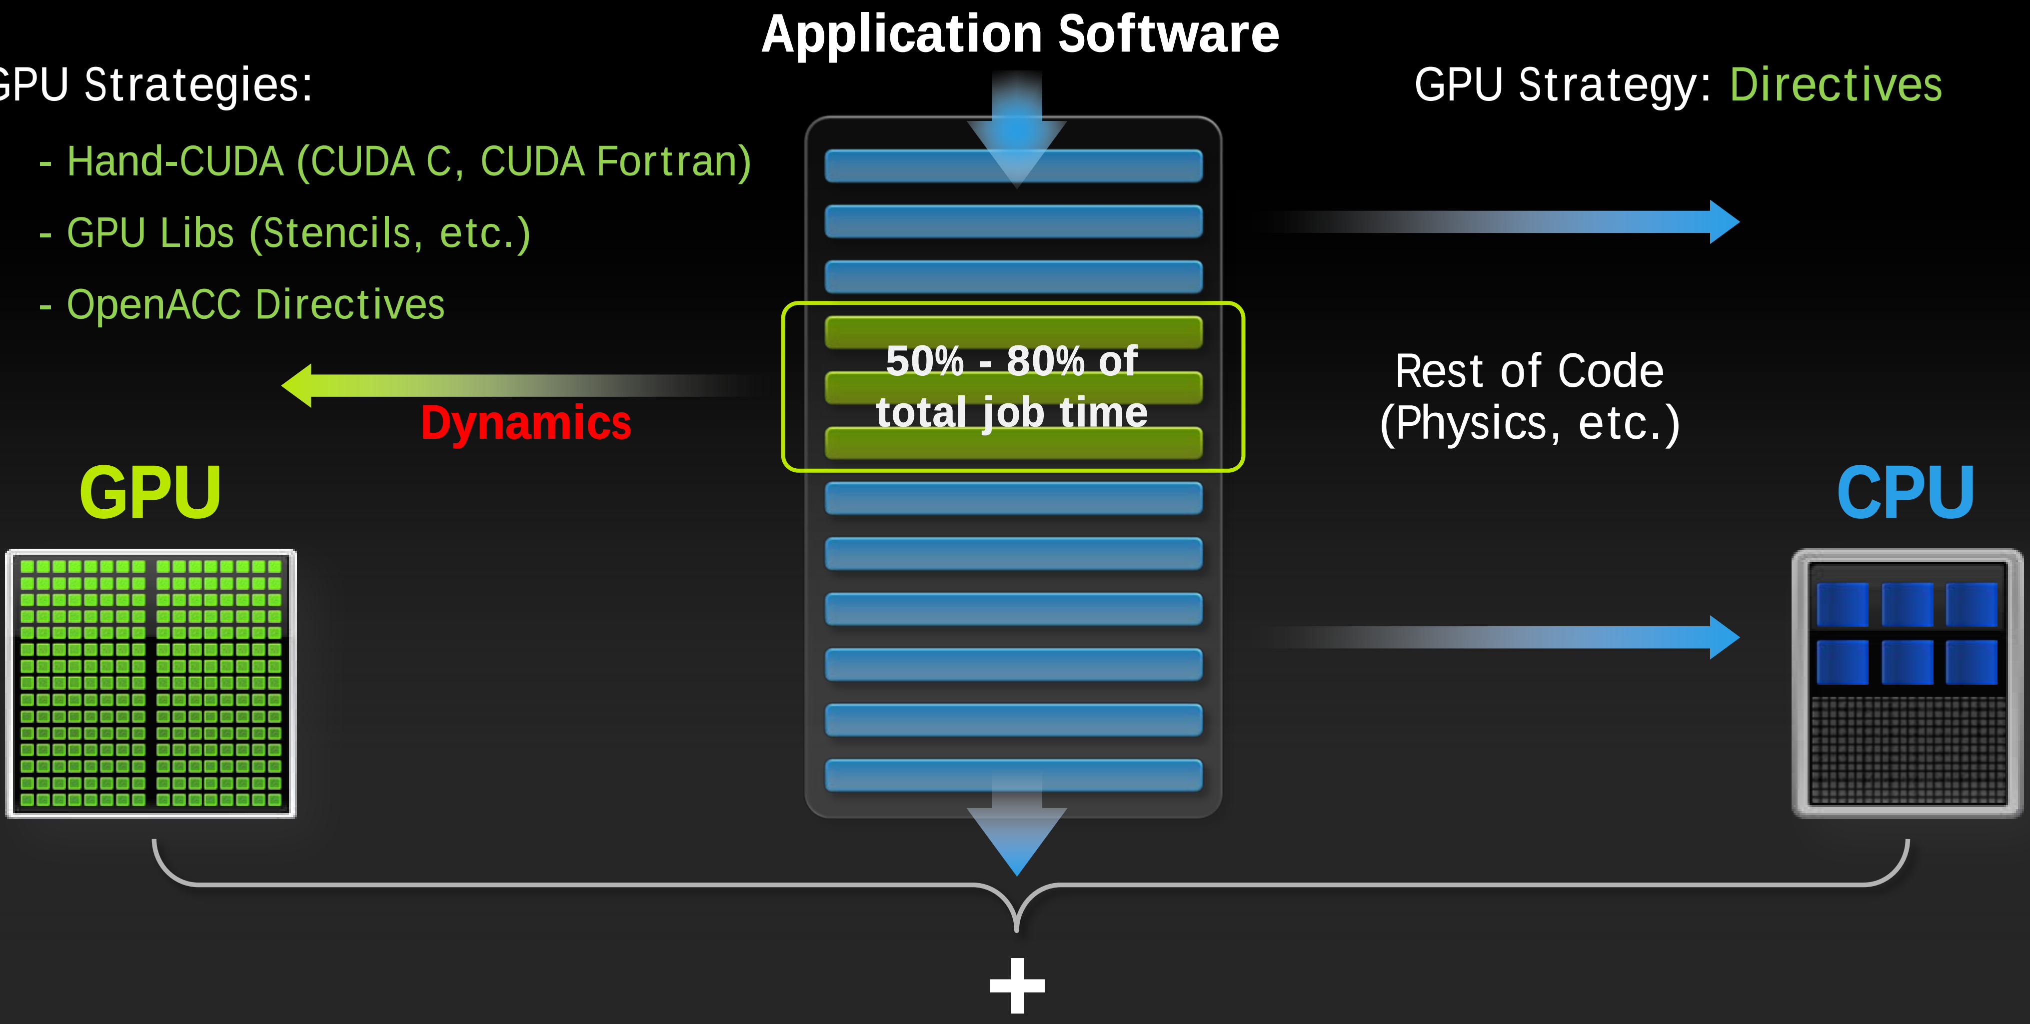

US DOE ORNL and CAM-SE Atmospheric Model

Programming weather, climate, and earth-system models on heterogeneous multi-core platforms

September 12-13, 2012 at the National Center for Atmospheric Research in Boulder, Colorado

Porting the Community Atmospheric Model - Spectral Element Code to Utilize GPU Accelerators

- Dr. Matt Norman, Oak Ridge National Laboratory

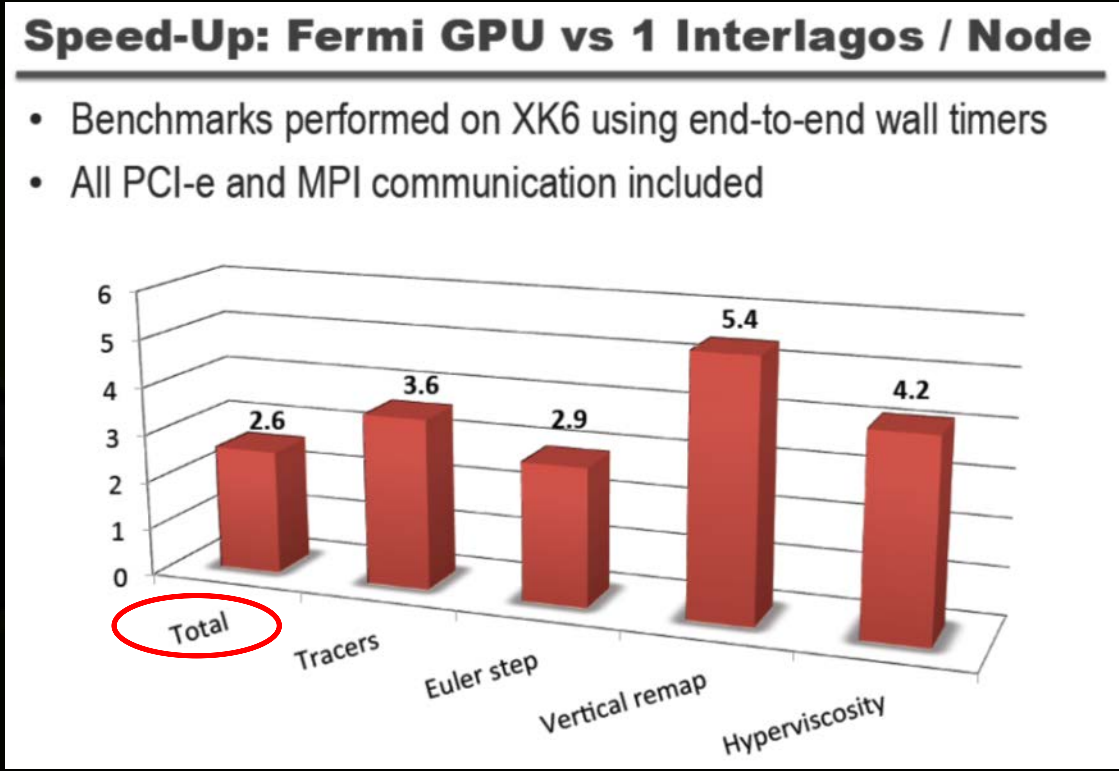

CAM-SE Dynamics

- HOMME spectral element dycore on cubed sphere
-
- CUDA Fortran today OpenACC option moving forward
- 2.6x speedup for 16 cores

Source: http://data1.gfdl.noaa.gov/multi-core/

CAM-SE a Critical Application for Titan at ORNL

World's Largest Open Science Computing Research Facility

14,592 NVIDIA Tesla GPUs

20+ PetaFlops

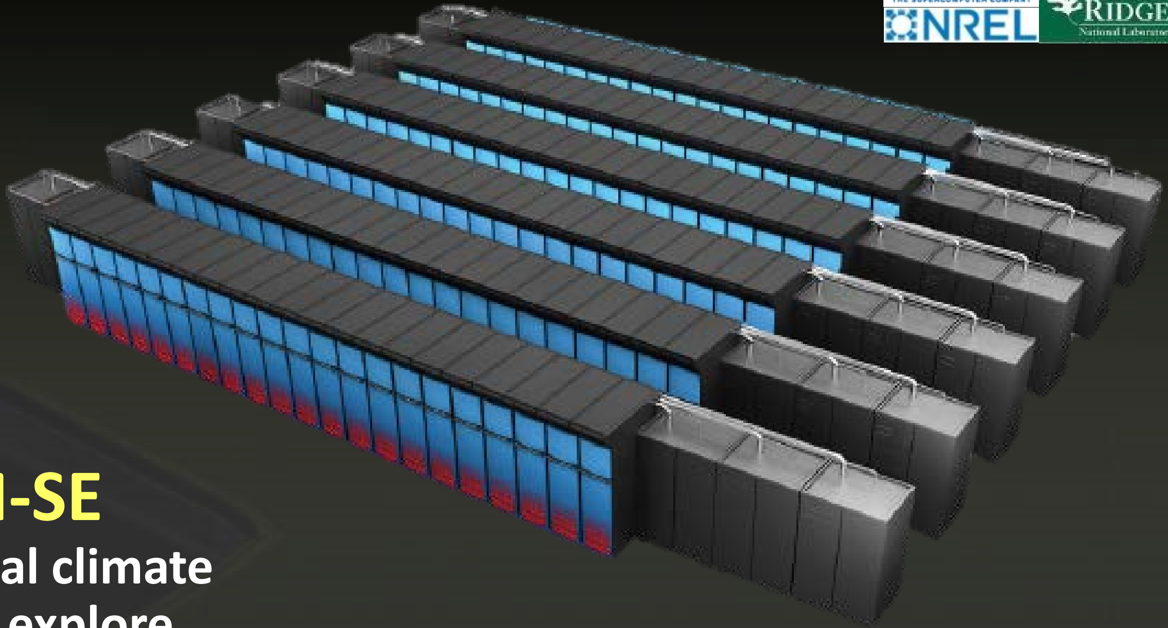

CAM-SE

Model global climate change & explore mitigation strategies

COSMO HP2C Project for GPU Development

- COSMO limited-area NWP/climate model - http://www.cosmo-model.org/
- Used by 7 weather services and O(50) universities and research institutes

Physics

- Large group of developers
- Plug-in code from other models
- Less memory bandwidth bound
- Simpler stencils (K-dependencies)
- 20% of runtime

→ Keep source code (Fortran)
→ GPU port with directives (OpenACC)

Dynamics

- Small group of developers
- Memory bandwidth bound
- Complex stencils (IJK-dependencies)
- 60% of runtime

→ Aggressive rewrite in C++
→ Development of a stencil library
→ Still single source code multiple library back-ends for x86 / GPU

Source: http://data1.gfdl.noaa.gov/multi-core/

GPU Motivation: Cost of Higher Resolution

Resolution is of key importance to increase simulation quality

2x resolution ≈ 10x computational cost

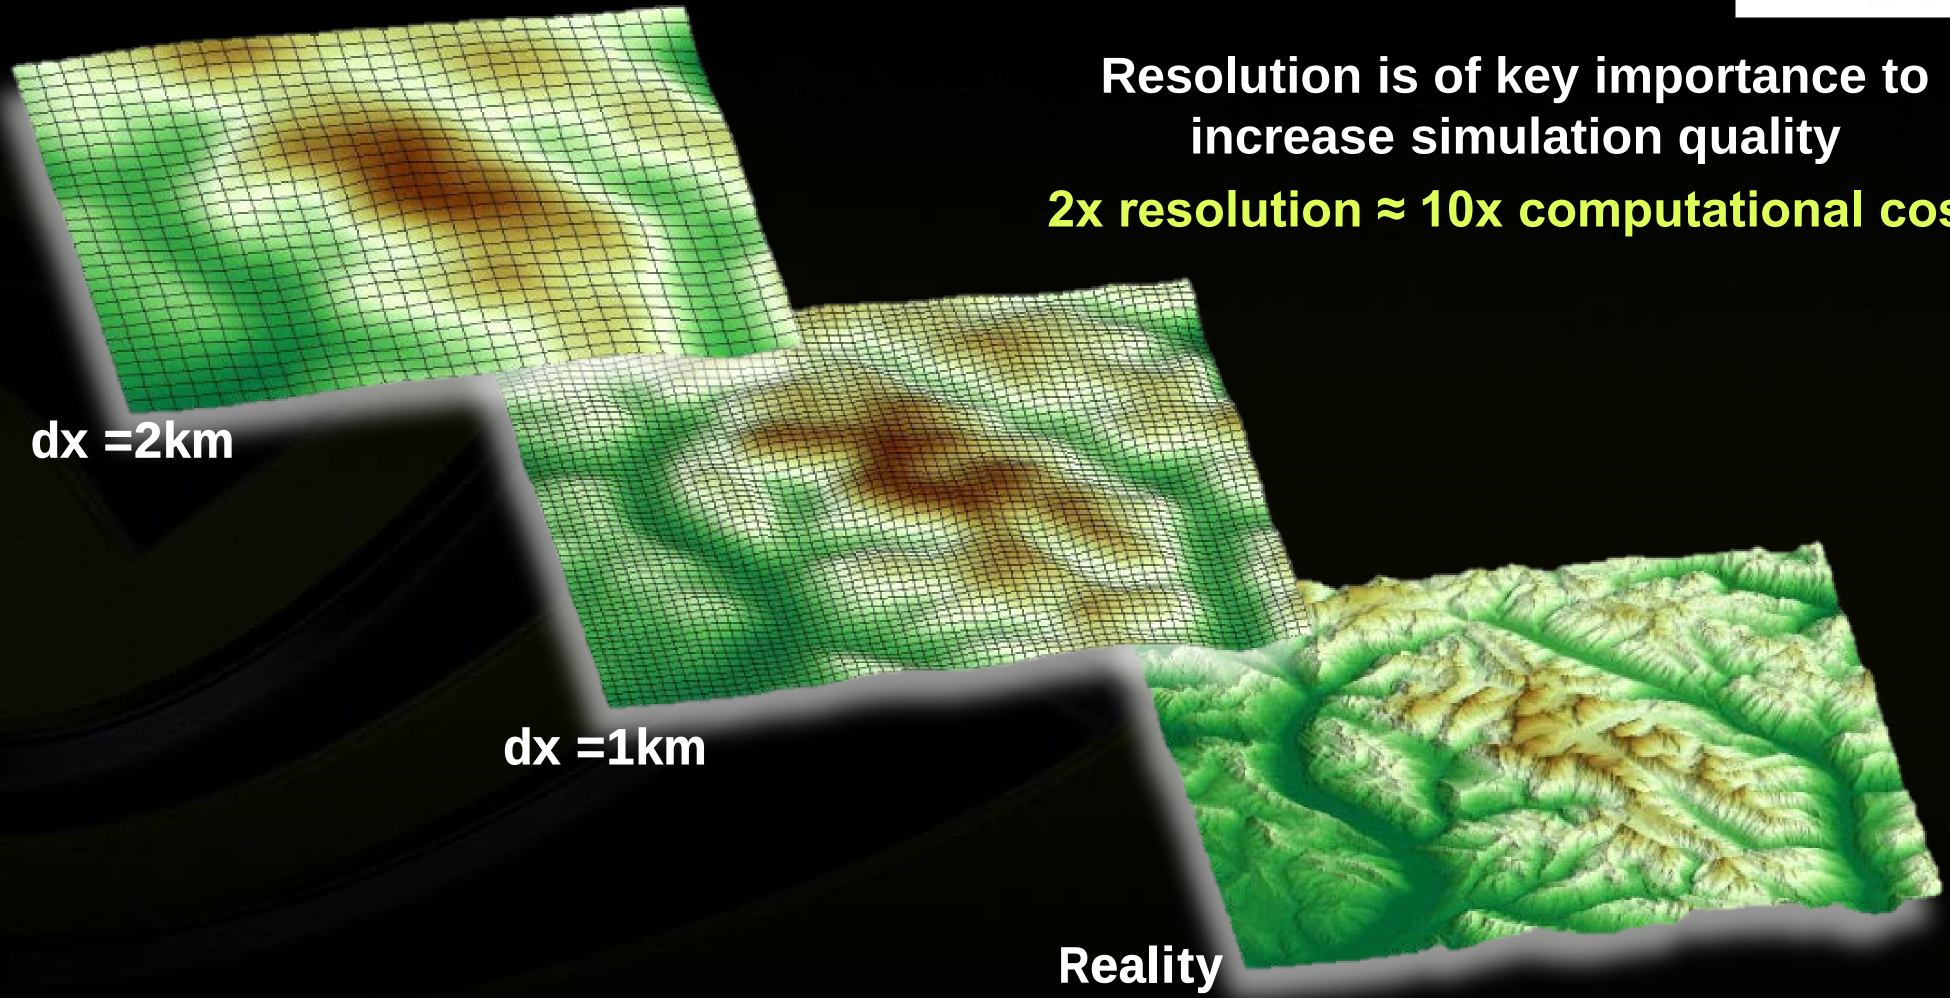

GPU Status of WRF Developments

- Several non-trunk efforts at various stages
  - Dynamics and some physics by Thomas Nipen at UBC – source at NVIDIA
  - KernelGen project: www.kernelgen.org update at NCAR 2012 workshop
  - Cray and OpenACC (Pete Johnsen) with results at 2012 NCAR workshop
  - C-DAC and HPC-FTE group working with NVIDIA India (Priyanka)
  - Shortwave radiation model by NV software group and PGI (G. Ruetsch)
  - Physics kernels by John Michalakes: www.mmm.ucar.edu/wrf/WG2/GPU/
  - NIM to include WRF physics using PGI and/or HMPP, OpenACC

- Trunk effort in release v3.2 from 2009
  - WSM5 physics model (15% - 25%)

UBC Developments of WRF

- **Dynamics and some physics by Thomas Nipen at UBC, with John Michalakes ~2010**

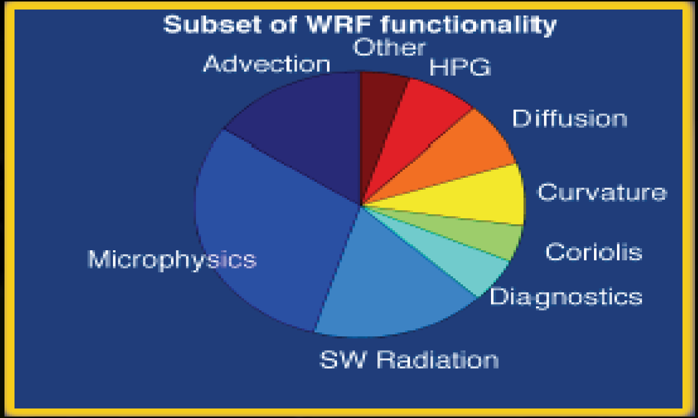

Subset of WRF model:
- Various dynamics components
- Microphysics (Kessler)
- Shortwave radiation (Dudhia)

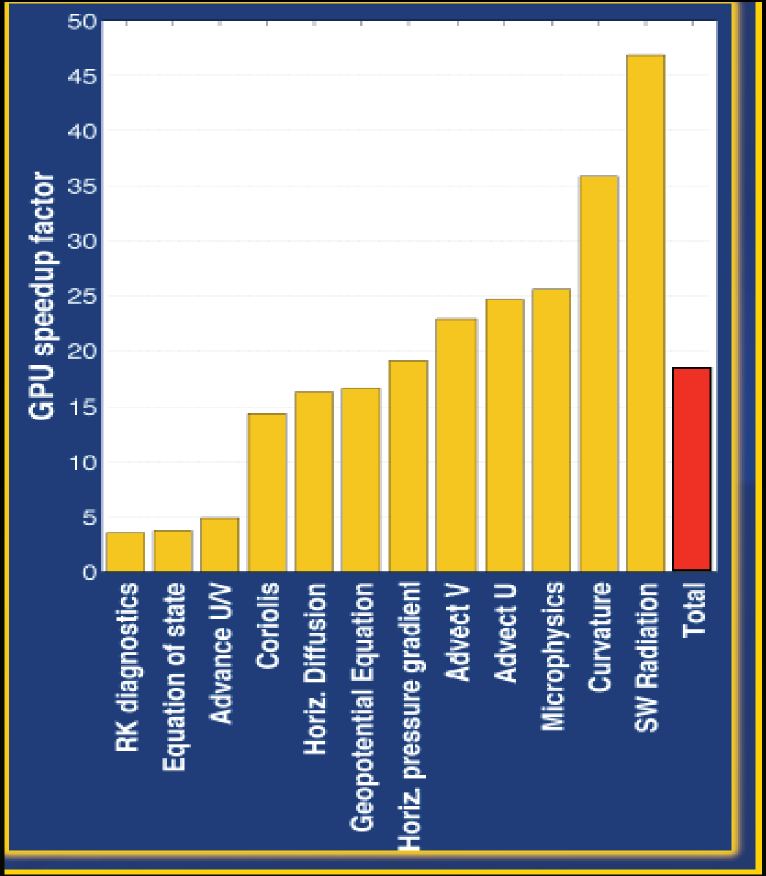

Simulation on:
- 189 x 150 x 27 domain
- CPU: 2.4GHz Opteron*
- GPU: GeForce 9800 GX2

*using one core

Source: https://ams.confex.com/ams/91Annual/webprogram/Paper181202.html

Developments using Directives and OpenACC

Model	Domain	Collaborators
NEMO	Ocean Model	NVIDIA
WRF	NWP/Climate	Cray, NVIDIA
COSMO	NWP/Climate	CSCS, SCS, NVIDIA
GEOS-5	Climate	NASA GSFC, PGI
NIM and FIM	NWP/Climate	NOAA, PGI, CAPS, Cray, NVIDIA

OpenACC
DIRECTIVES FOR ACCELERATORS

www.openacc-standard.org

NEMO Acceleration Using OpenACC

Background

- NEMO ocean modeling framework: http://www.nemo-ocean.eu/
- Used by 240 projects in 27 countries (14 in Europe, 13 elsewhere)

Approach

- Project based on NEMO 3.4, use of PGI Fortran compiler 12.9 preview
- Flat profile, 1st routine is 6%, many routines to accelerate for overall benefit
- OpenACC “present” clause keeps data on the device between subroutine calls
- Directives for 41 routines: rearranged loops in 12, temporary arrays in 13
- Other changes for improved MPI communication, other miscellaneous

NEMO Acceleration Using OpenACC

GYRE_50 Configuration, I/O disabled, OpenACC 1.0: **Speedup ~5x**

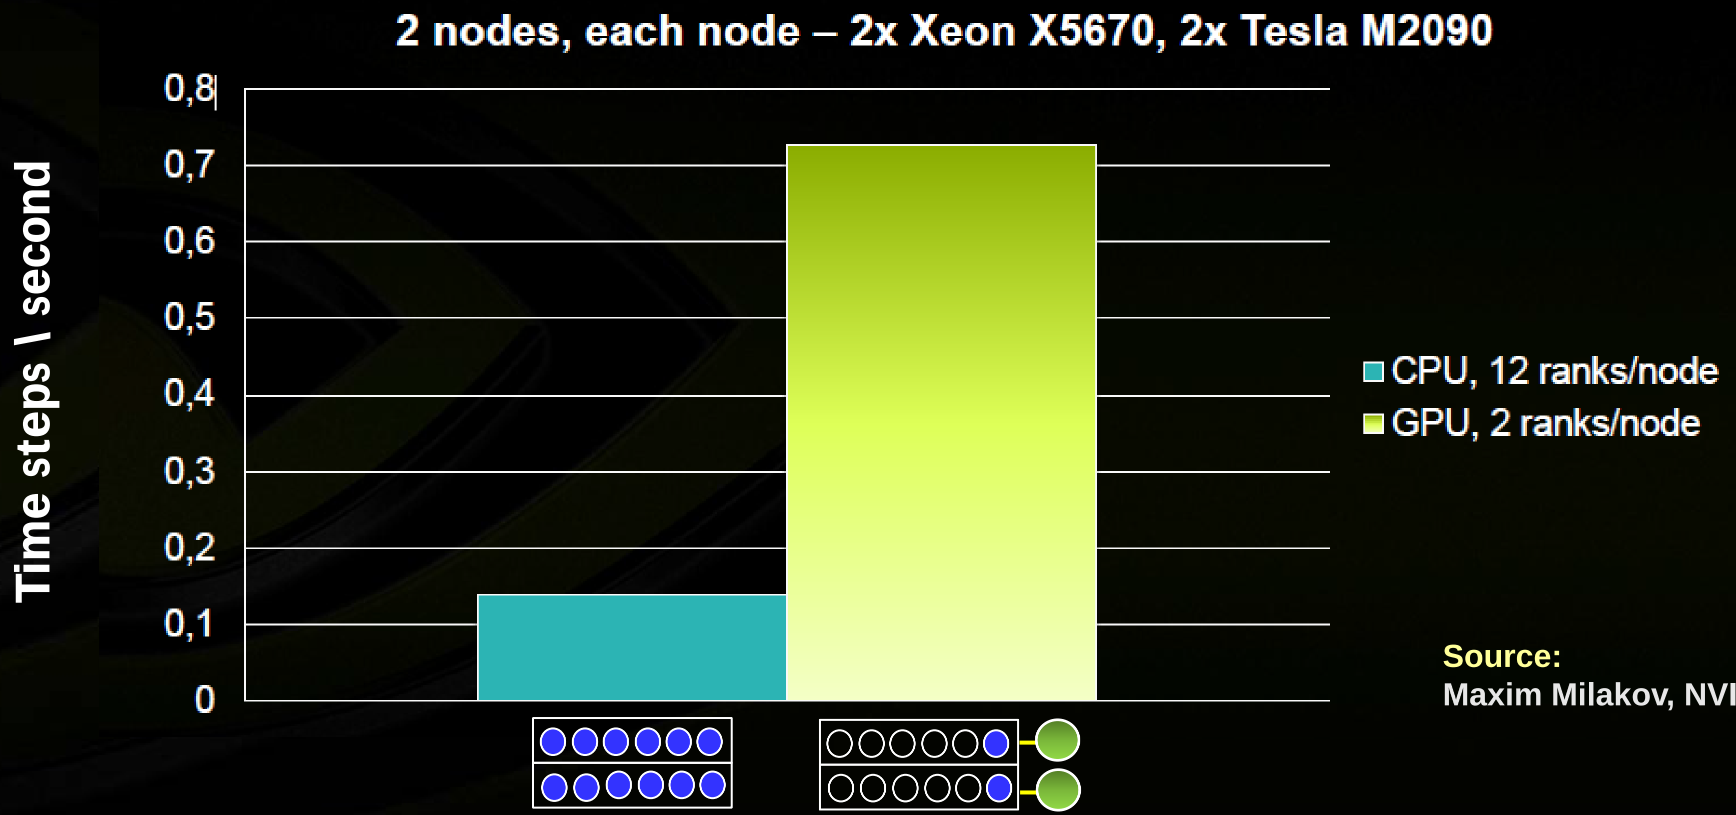

Source:
Maxim Milakov, NVIDIA

Agenda: Kepler GPU Architecture and Benefits to Earth System Modeling

- Introduction of GPUs in HPC
- Review of GPU Progress
- Kepler Architecture Benefits

NVIDIA CUDA GPU Roadmap

NVIDIA

NVIDIA Roadmap Characteristics:

* Increasing number of more flexible cores
* Larger and faster memories (6 GB today)
* Enhanced programming and standards
* Tighter integration with systems

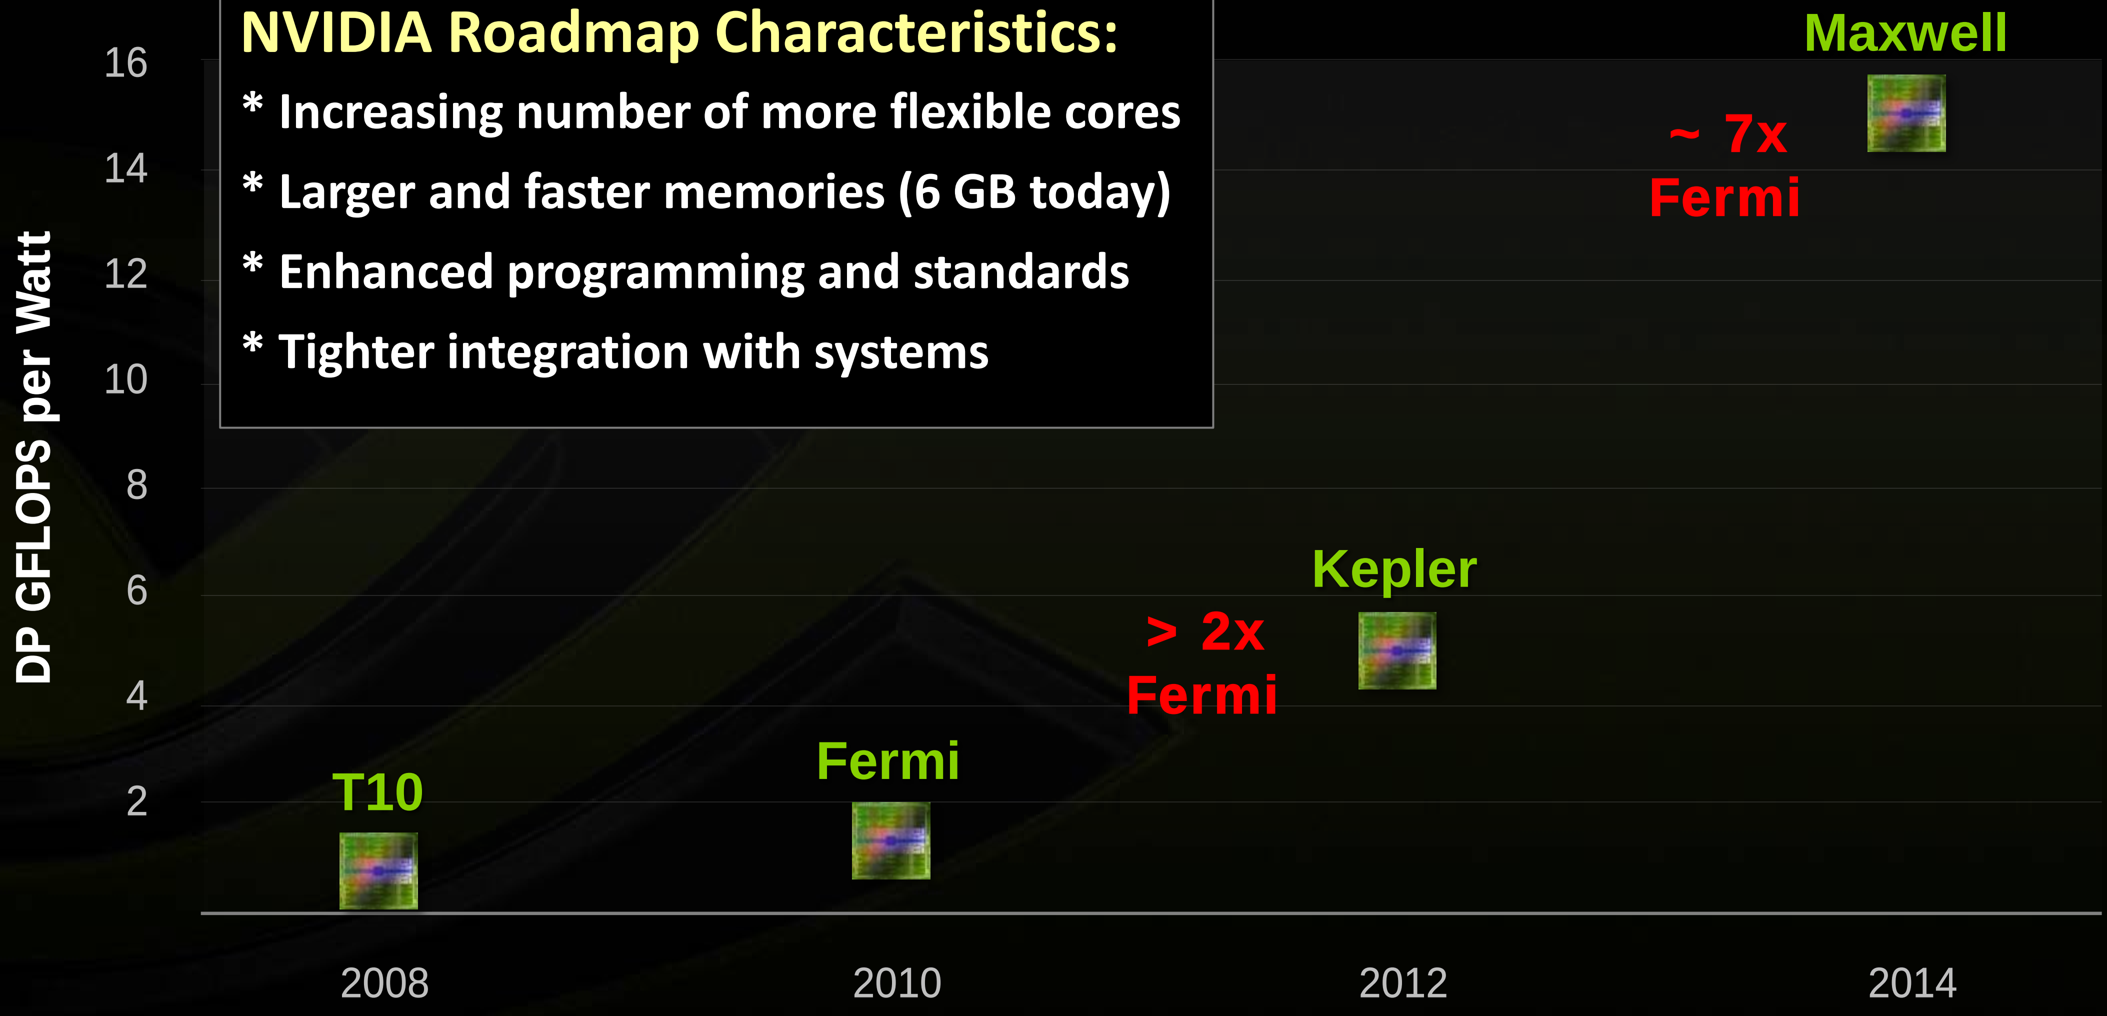

New GPU Architecture Kepler vs. Fermi

Features	Tesla M2075	Tesla K10	Tesla K20
GPU Architecture	Fermi GF100 	Kepler GK104 	Kepler GK110 
Peak DP performance	515 Gigaflops	190 Gigaflops (95 Gflops per GPU)	> 1000 Gigaflops
Peak SP performance	1030 Gigaflops	4577 Gigaflops (2.3 Tflops per GPU)	> 1000 Gigaflops
Memory bandwidth	150 GBytes/sec	320 GBytes/sec (160 per GPU)	> 200 GBytes/sec
Memory size (GDDR5)	6 Giga Bytes	8 Giga Bytes (4 GB per GPU)	6 Giga Bytes
CUDA cores	448	3072 (1536 per GPU)	TBA
Available in Servers	Since 2010	Today	Late 2012

Kepler Architecture and CUDA 5 Highlights

- **SMX**
  - **New instructions for programmability**
- **CUDA Dynamic Parallelism (CDP)**
  - **More responsibility for GPU**
- **Hyper-Q**
  - **Improved GPU utilization**

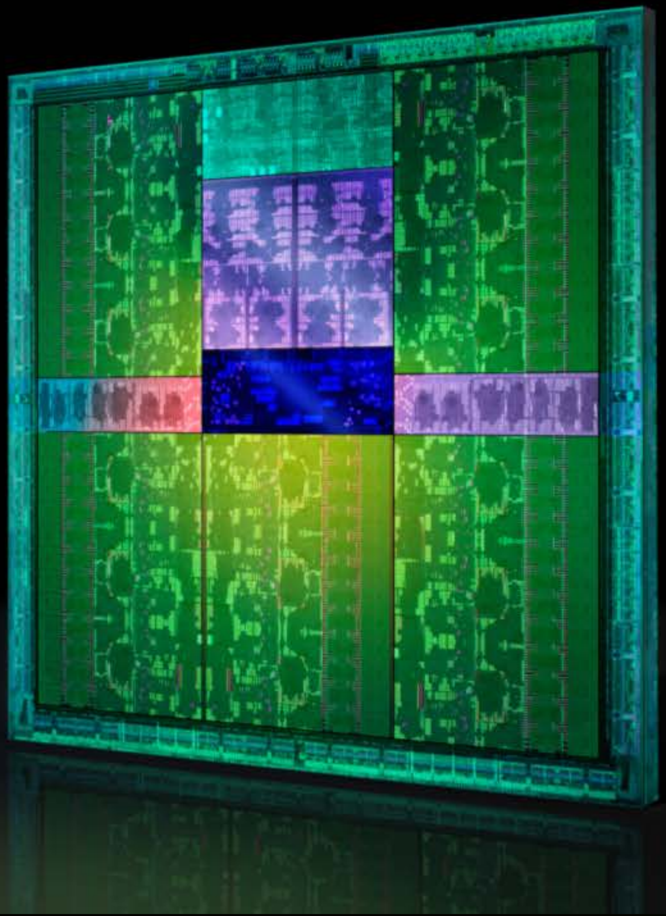

Streaming Multiprocessor SMX

SMX Resource	Kepler K20 vs. Fermi
Floating point throughput	2-3x
Max Blocks per SMX	2x
Max Threads per SMX	1.3x
Register File Bandwidth	2x
Register File Capacity	2x
Shared Memory Bandwidth	2x
Shared Memory Capacity	1x

CUDA Dynamic Parallelism (CDP)

Ability to launch new grids directly from the GPU

- Dynamically
- Simultaneously
- Independently

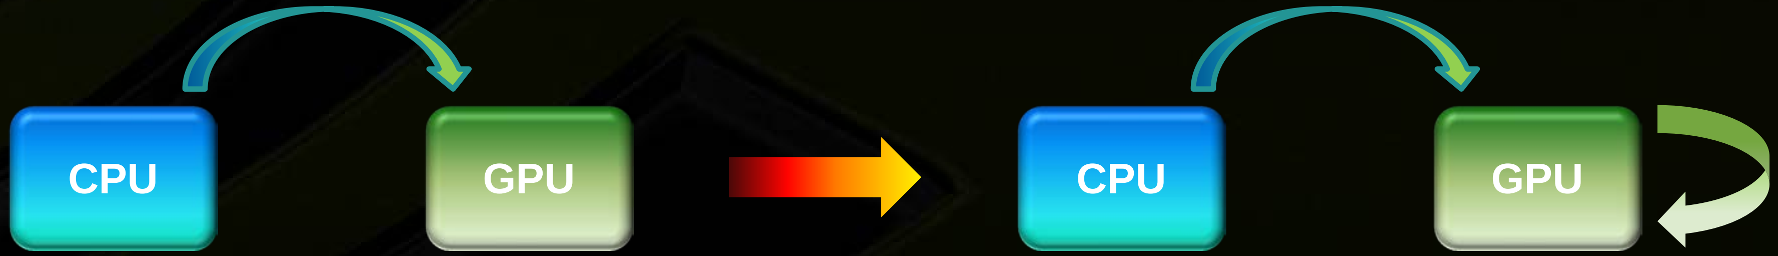

Fermi: Only CPU can generate GPU work

Kepler: GPU can generate work for itself

CUDA Dynamic Parallelism (CDP)

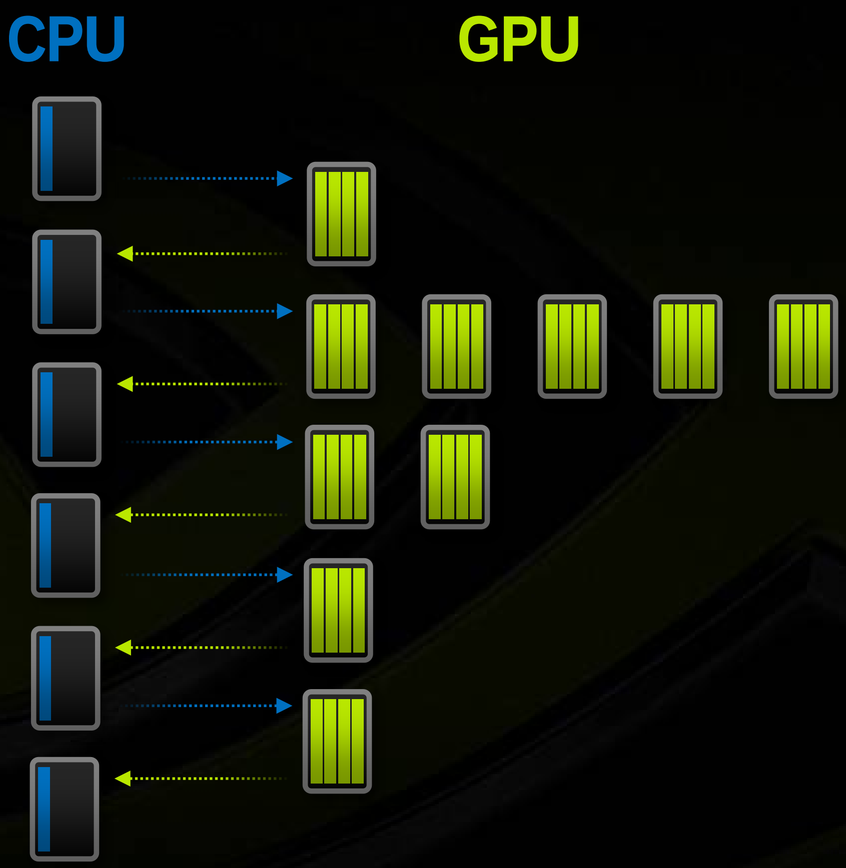

GPU as Co-Processor

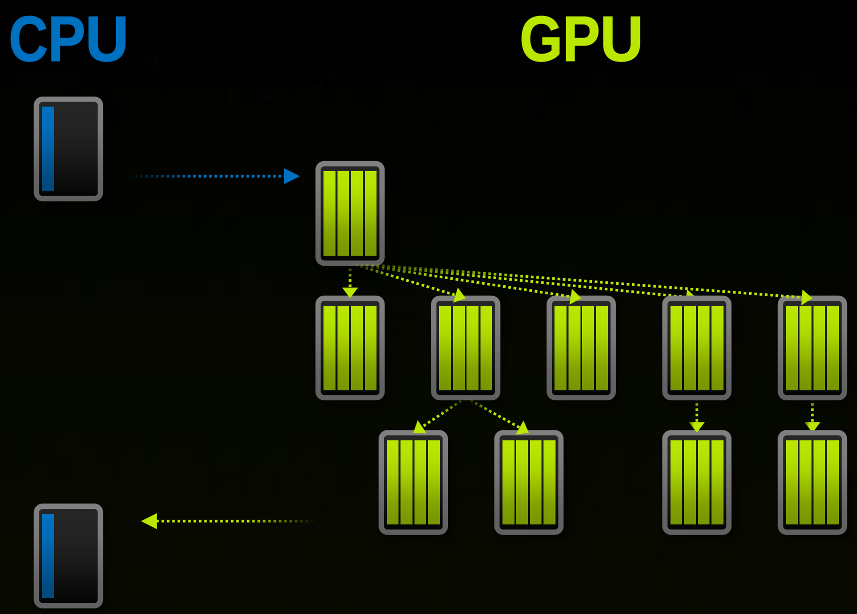

Autonomous, Dynamic Parallelism

Hyper-Q Improves Efficiency for MPI Parallel

CPU Cores Simultaneously Run Tasks on Kepler

FERMI
1 MPI Task at a Time

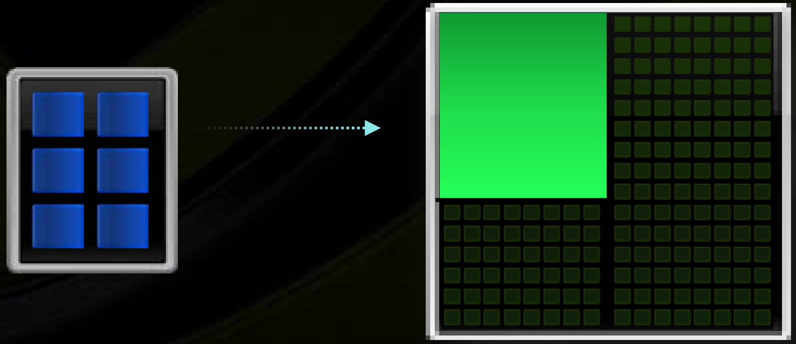

KEPLER
32 Simultaneous MPI Tasks

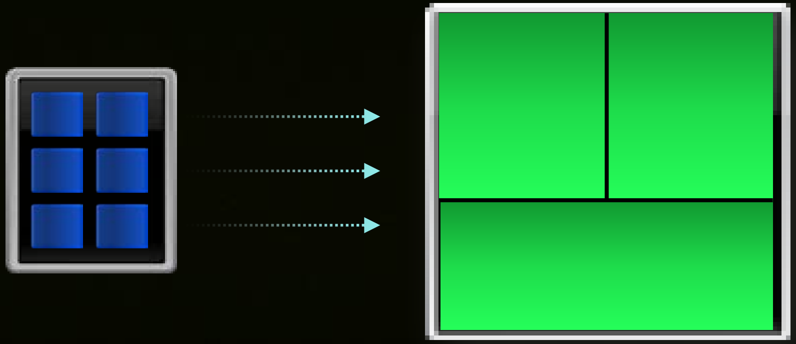

Hyper-Q Improves Efficiency for MPI Parallel

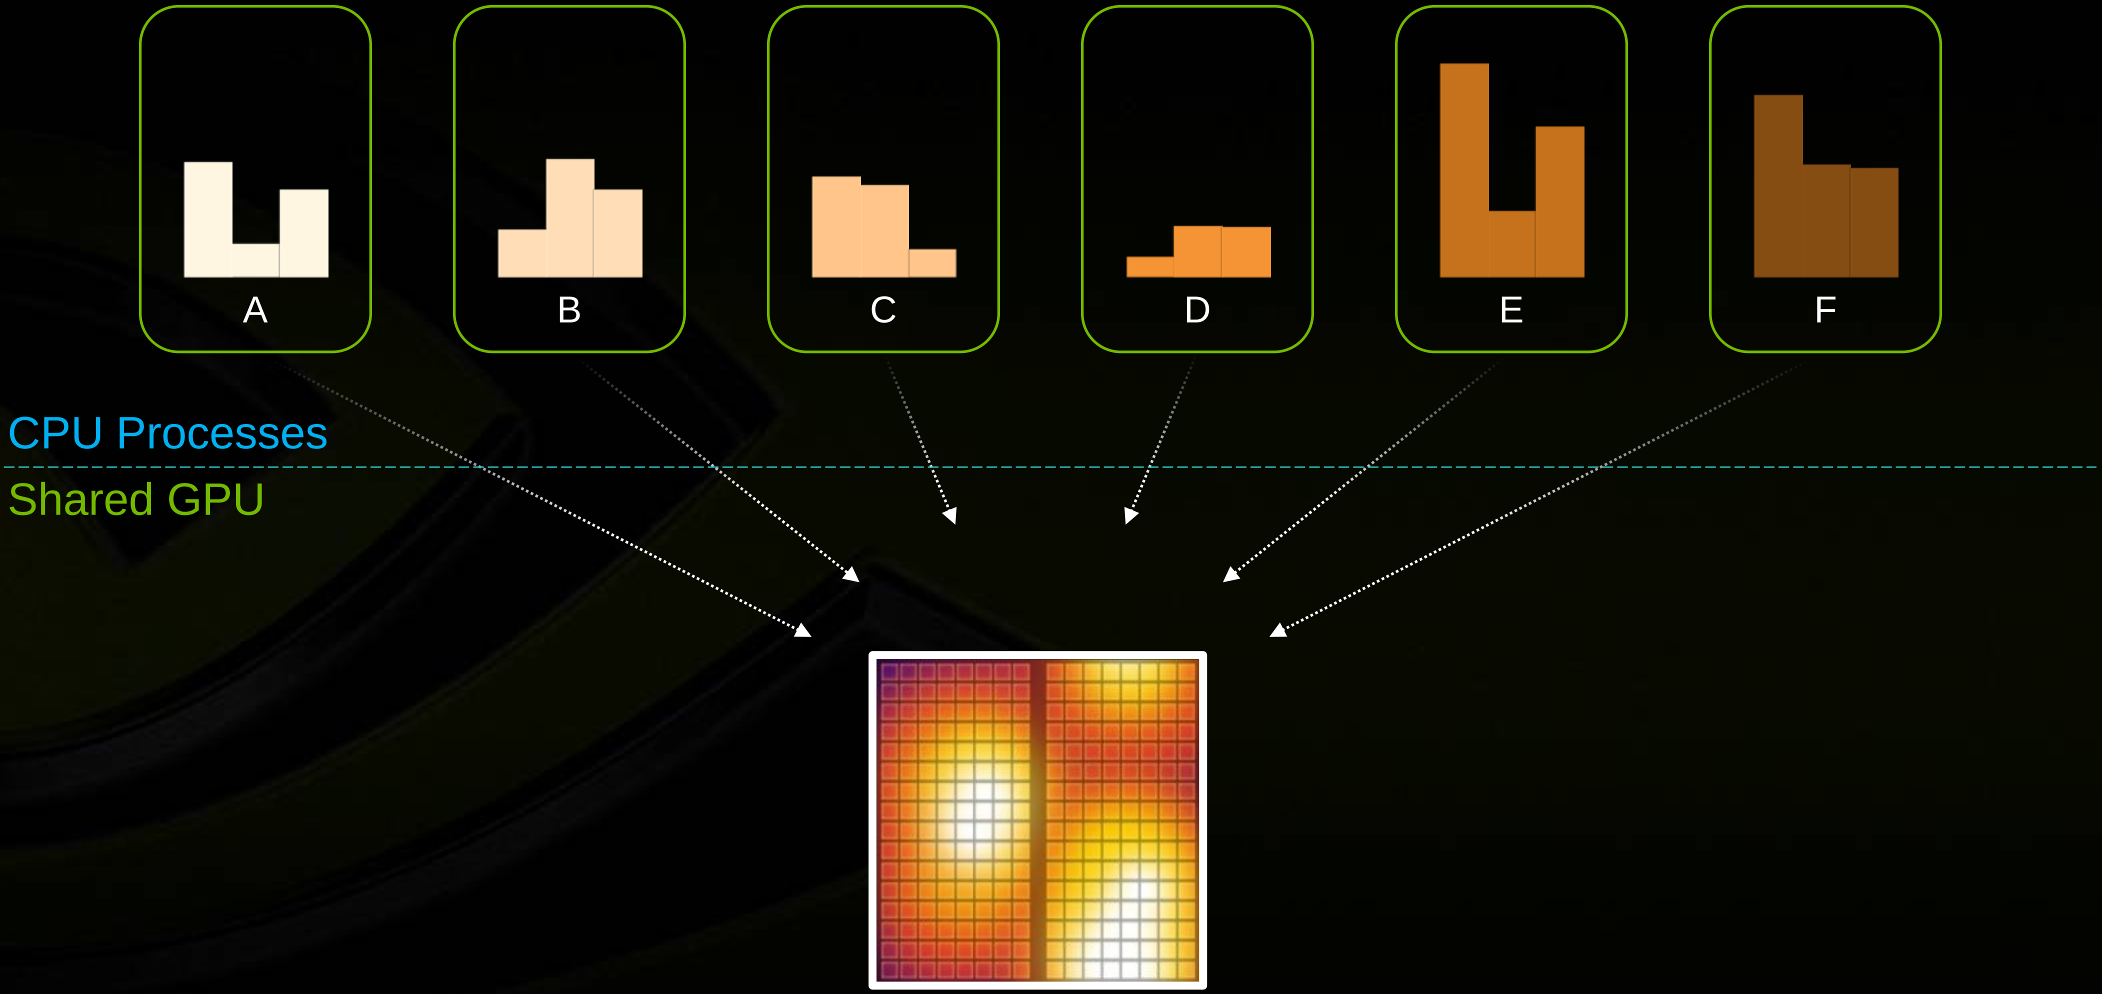

Hyper-Q Improves Efficiency for MPI Parallel

- **Streams from multiple CPU processes can execute concurrently**

- **Use as many MPI ranks as in CPU-only case**

=> smaller impact of CPU work

- **Particularly interesting for strong scaling**

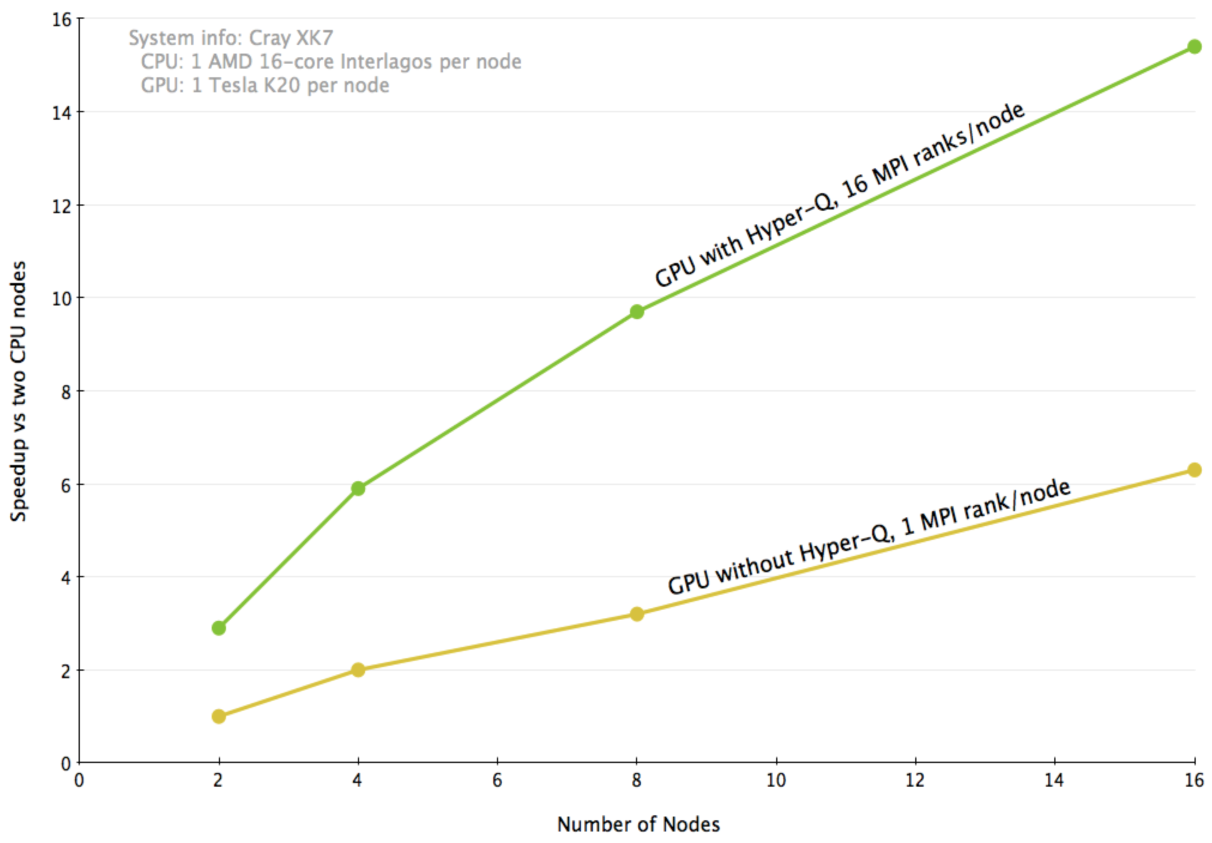

Kepler Enables Full NVIDIA GPUDirect™

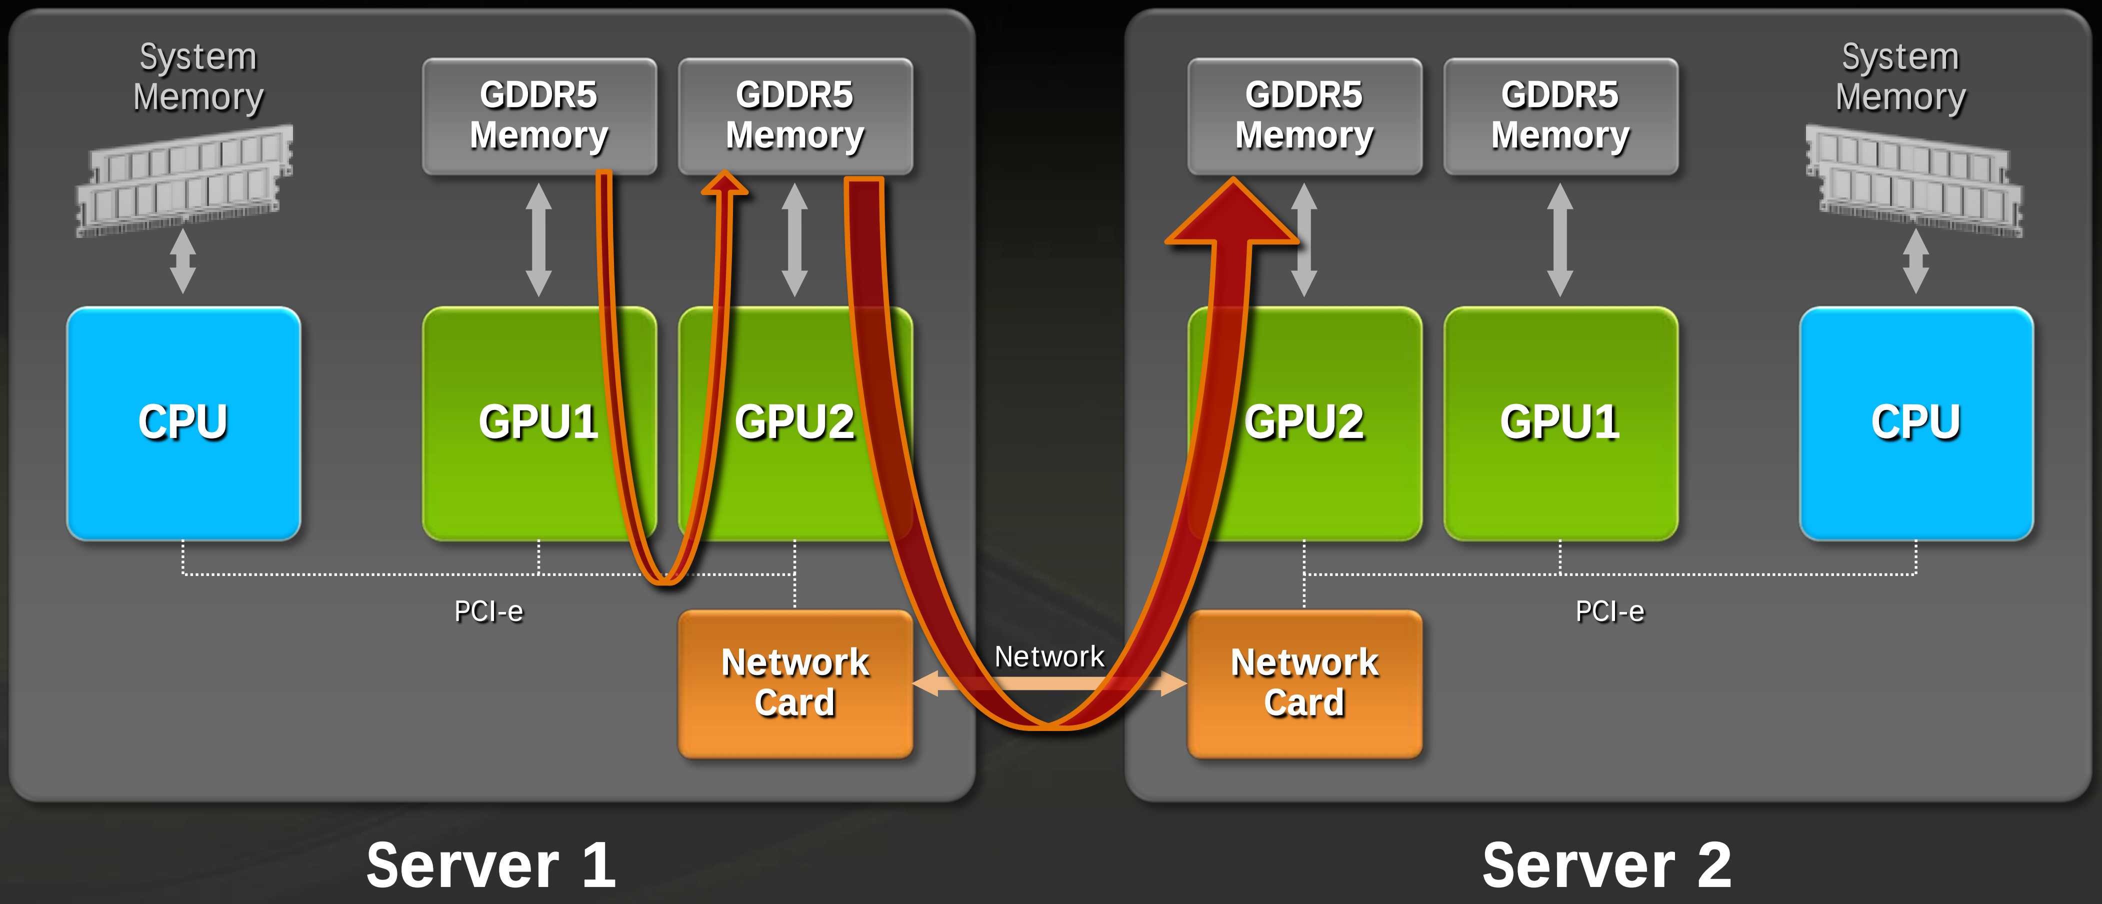

Remote Visualization with NVIDIA VGX

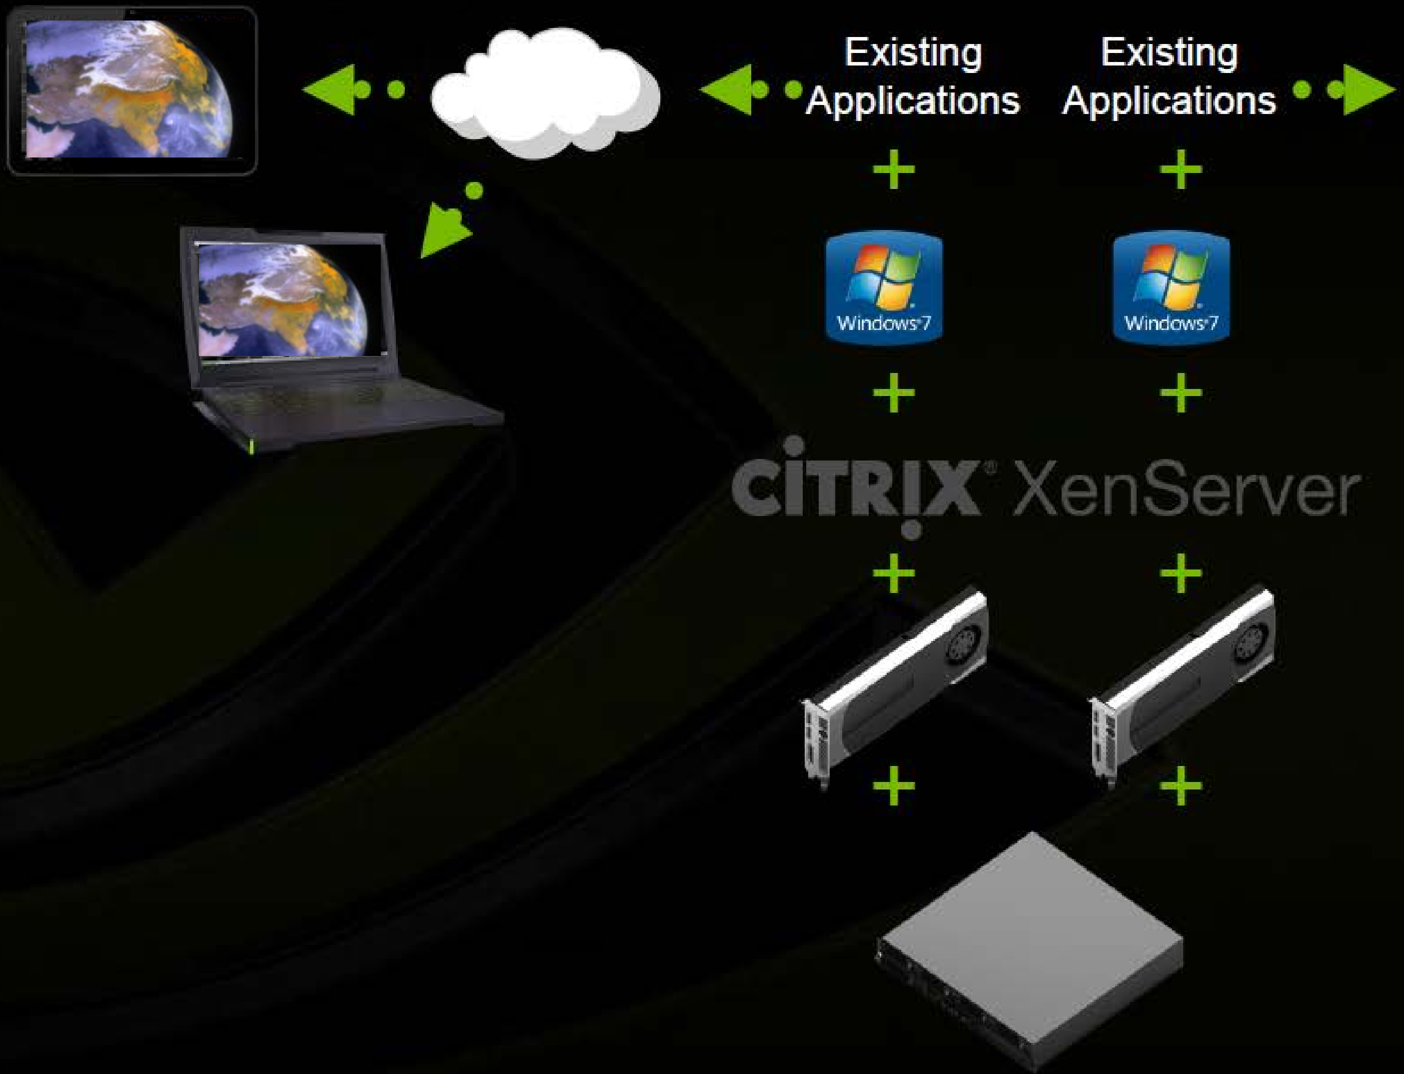

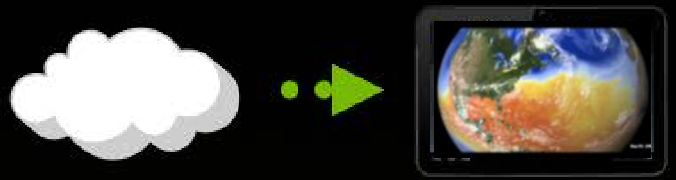

NVIDIA Monterey Remote Technology
Delivers remote graphics with local look and feel

Citrix XenServer 6 with NVIDIA Multi-OS Technology
Enables high-performance Quadro graphics on a virtual machine

NVIDIA Quadro Graphics
Brings NVIDIA® Quadro® and Tesla™ M2070Q performance to the data center

Summary

- Several NWP/Climate/Ocean Models Support NVIDIA GPUs
  - Mostly dynamics today but focus on full model implementations
  - New Kepler architecture offers several benefits to ES modeling
- NVIDIA Investments in Collaborations with Key Organizations
  - ES modeling community, research orgs, system vendors, OpenACC
- Learn More About NVIDIA HPC Solutions
  - More at: www.nvidia.com/tesla
  - White paper on Kepler architecture:
    www.nvidia.com/content/PDF/kepler/NVIDIA-Kepler-GK110-Architecture-Whitepaper.pdf
  - Want to investigate GPUs, please contact sposey@nvidia.com

Thank You, Questions?

Stan Posey | HPC Industry and Applications Development | sposey@nvidia.com