GPU Considerations of ESMs for the Pre-Exascale HPC Phase

Stan Posey; sposey@nvidia.com; NVIDIA, Santa Clara, CA, USA
HPC Program Manager, ESM and CFD Solutions

Agenda: Application Readiness for the Pre-Exascale Phase

- NVIDIA Update and HPC Trends
- Developments in Pre-Exascale HPC
- ESM Requirements and GPU Progress

NVIDIA Core Technologies and Markets

Company Revenue of ~$5B USD; ~9,000 Employees; HPC Growing > 30% CAGR

PC | DATA CENTER | MOBILE

Tesla®

Overview of Tesla GPU Acceleration Concept

NVIDIA

Optimization of Serial + Parallel Execution

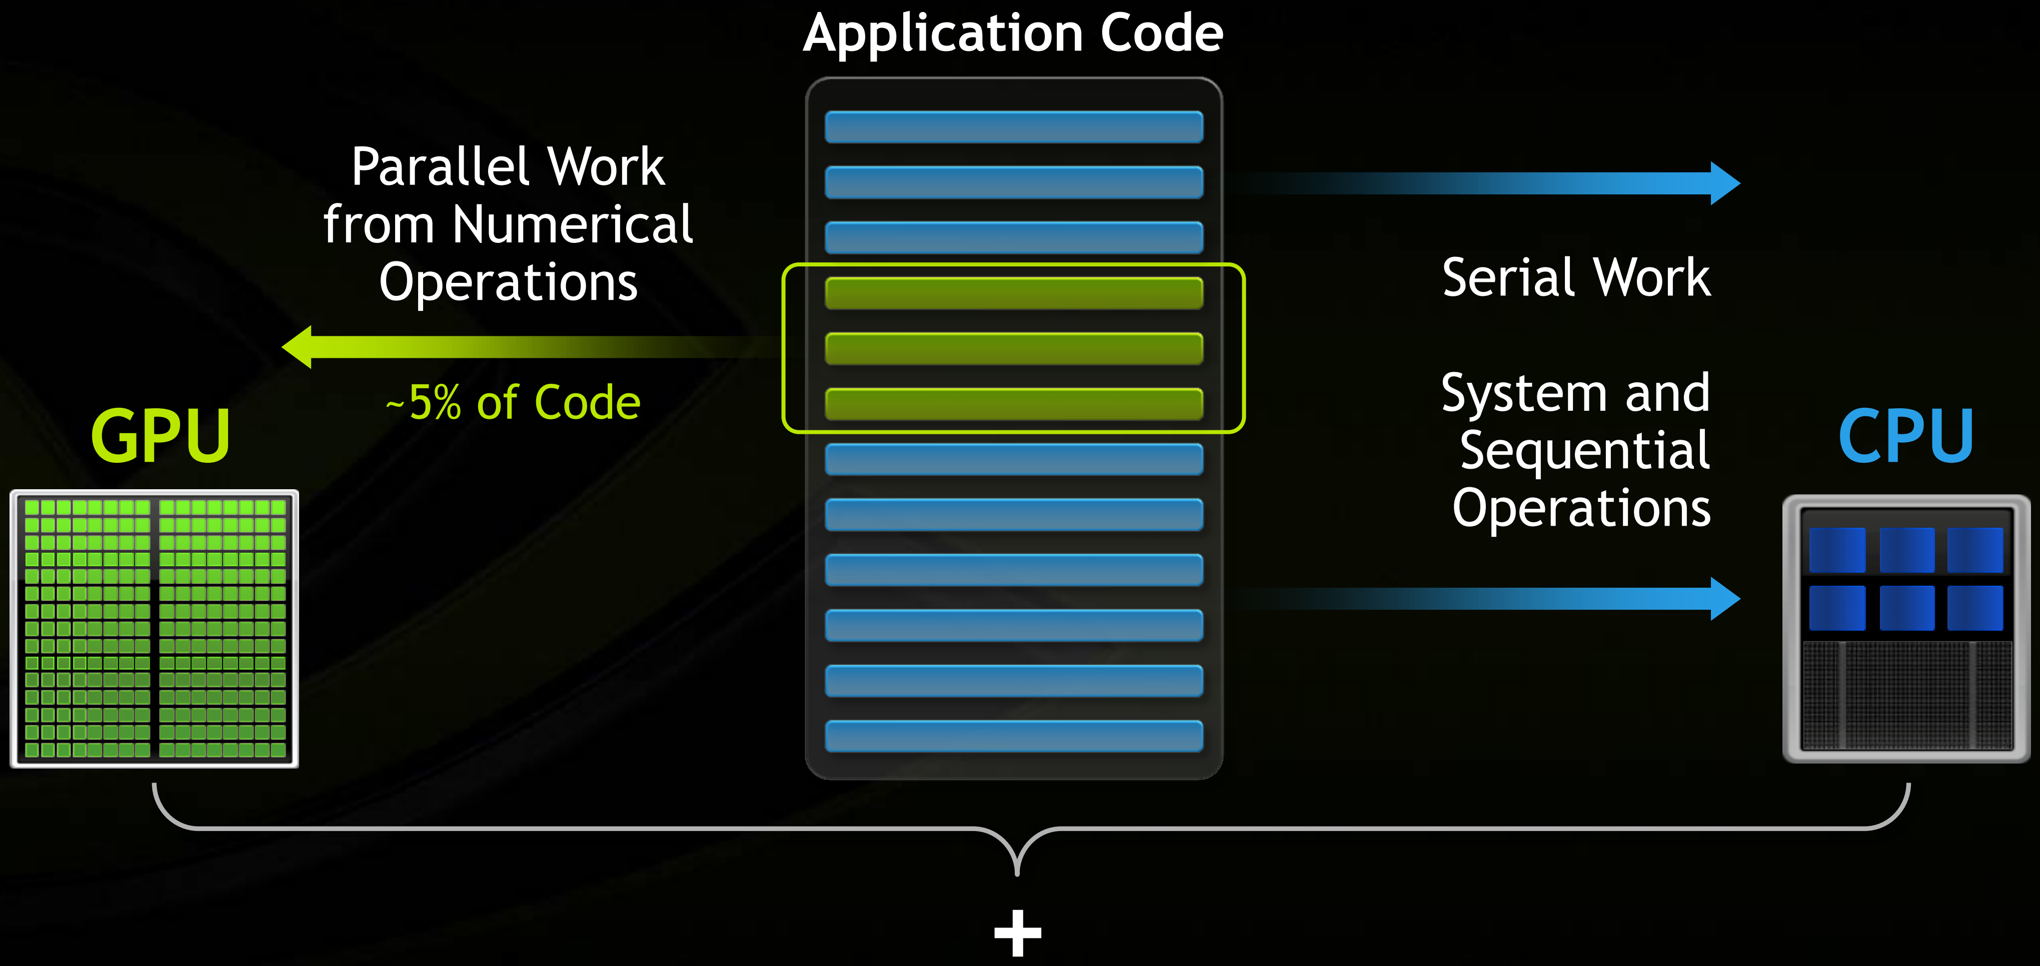

GPU Motivation (II): Power Limiting Trends

Challenges of Getting ECMWF's Weather Forecast Model (IFS) to the Exascale
– G. Mozdzynski, ECMWF 16th HPC Workshop

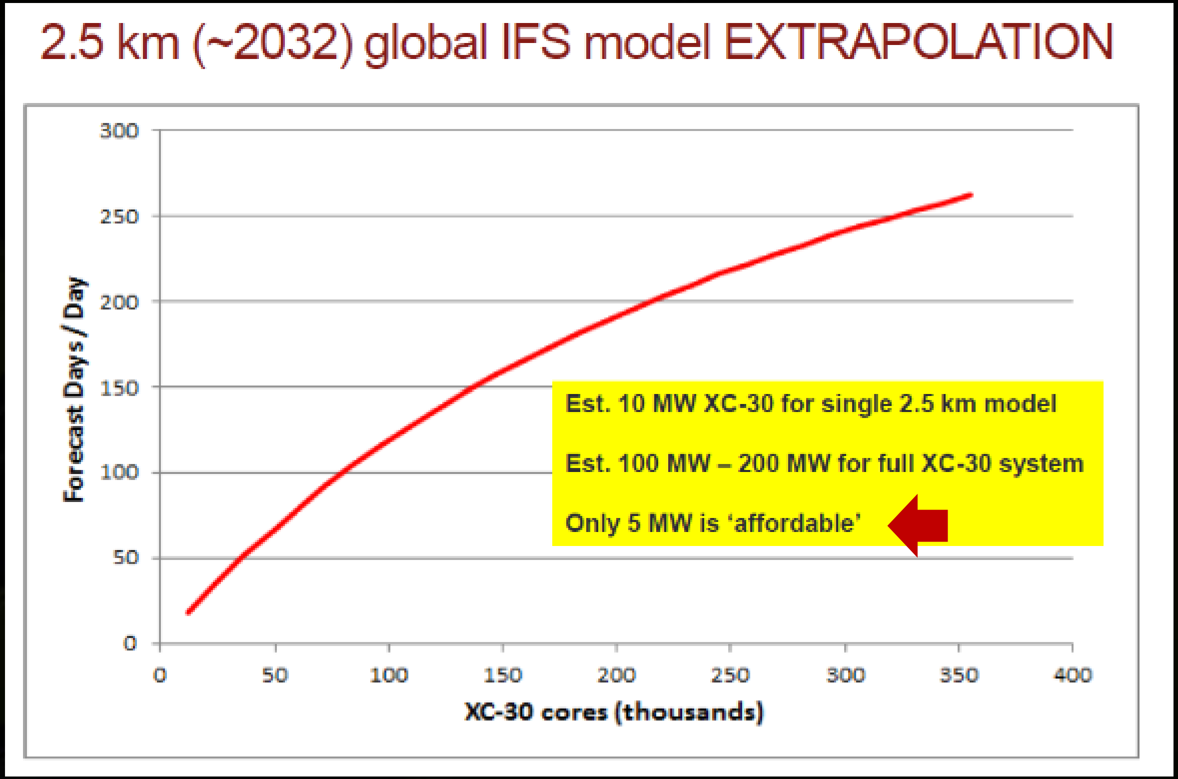

"**Conclusion**: Current mathematics, algorithms and hardware will not be sufficient (by far)"

GPU Motivation (III): ES Model Trends

- **Higher grid resolution with manageable compute and energy costs**
  - Global atmosphere models from 10 km today to cloud-resolving scales of 1-3 km

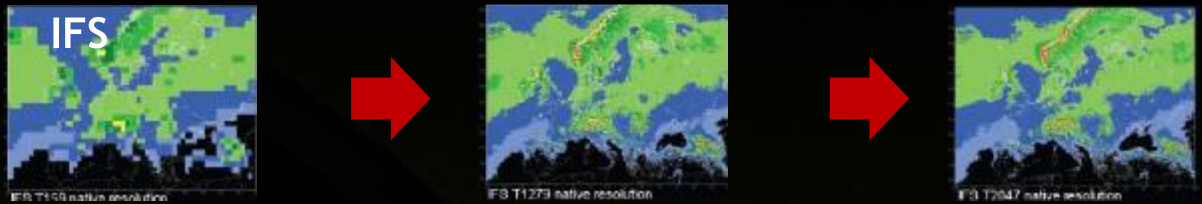

128 km → 16 km → 10 km → ? 3 km

Source: Project Athena – http://www.wxmaps.org/athena/home/

- **Increase in ensemble use and ensemble members to manage uncertainty**

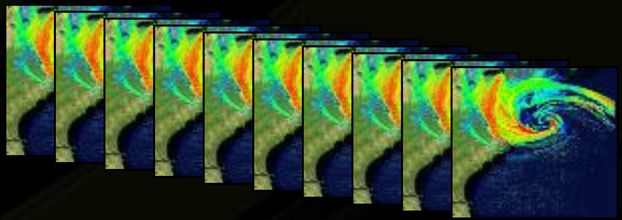

Number of Jobs 10x - 50x

- **Fewer model approximations (NH), more features (physics, chemistry, etc.)**

Accelerator technology identified as a cost-effective and practical approach to future NWP requirements

Tesla GPU Progression During Recent Years

	2012 (Fermi) **M2075**	2014 (Kepler) **K20X**	2014 (Kepler) **K40**	2014 (Kepler) **K80**	**Kepler / Fermi**
Peak SP Peak SGEMM	1.03 TF	3.93 TF 2.95 TF	4.29 TF 3.22 TF	8.74 TF	**4x**
Peak DP Peak DGEMM	.515 TF	1.31 TF 1.22 TF	1.43 TF 1.33 TF	2.90 TF	**3x**
Memory size	6 GB	6 GB	**12 GB**	**24 GB (12 each)**	**2x**
Mem BW (ECC off)	150 GB/s	250 GB/s	**288 GB/s**	**480 GB/s (240 each)**	**2x**
Memory Clock		2.6 GHz	3.0 GHz	3.0 GHz	
PCIe Gen	Gen 2	Gen 2	Gen 3	Gen 3	**2x**
# of Cores	448	2688	2880	4992 (2496 each)	**5x**
Board Power	**235W**	**235W**	**235W**	**300W**	**0% – 28%**

Note: Tesla K80 specifications are shown as aggregate of two GPUs on a single board

GPU Technology a Mainstream HPC Platform

NVIDIA

Data Center Infrastructure

System Solutions
- CRAY
- DELL
- hp
- IBM
- . . .

Communication
- Mellanox Technologies
- MVAPICH
- Open MPI

Infrastructure Management
- Adaptive Computing
- Bright Computing
- IBM Platform Computing

Development

Programming Languages
- C/C++
- Fortran
- OpenACC
- python

Development Tools
- allinea DDT
- PGI
- VAMPIR

Software Solutions
- Kitware
- MATLAB
- Rogue Wave Software

GPU Accelerators
GPU Boost

Interconnect
GPU Direct
NVLink

System Management
NVML

Compiler Solutions
LLVM

Profile and Debug
CUPTI

Accelerated Libraries
cuBLAS

Enterprise Services Support & Maintenance

GPUs Applied Across Mainstream HPC Domains

NVIDIA

Oil & Gas

- Schlumberger
- BR Petrobras
- Eni
- Chevron
- Statoil

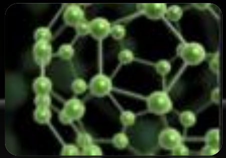

Higher Ed

- Harvard School of Engineering and Applied Sciences
- Stanford University
- Georgia Tech
- ETH
- University of Cambridge

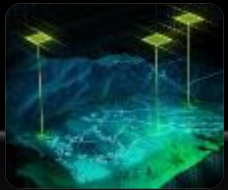

Government

- Air Force Research Laboratory
- Raytheon
- NASA
- Naval Research Laboratory

Supercomputing

- CSCS
- NCSA
- Tokyo Institute of Technology
- Oak Ridge National Laboratory
- Lawrence Livermore National Laboratory

Finance

- J.P.Morgan
- Barclays
- Standard Life
- BNP Paribas
- Murex

Consumer Web

- Baidu 百度
- salesforce
- Shazam
- amazon.com
- Yandex

GPUs Expanding from HPC to Data Analytics

NVIDIA

Attendees at NVIDIA GPU Technology Conference (GTC) 2014

- BOEING
- Chevron
- GE
- Harvard
- Stanford University
- NASA
- NCSA
- NOAA
- Oak Ridge National Laboratory
- Pacific Northwest National Laboratory
- SIEMENS
- Raytheon

- Baidu 百度
- Coupons.com
- Dropbox
- facebook
- Google
- NAVER
- NETFLIX
- Palantir
- Square
- Walmart
- YAHOO!
- Yandex

Broad Use of GPUs in Deep Learning

NVIDIA

HPC Drivers for Deep Learning → More Data

Better Models

GPUs

Early Adopters

Adobe
Image Analytics for Creative Cloud

Baidu
Speech/Image Recognition

flickr
Image Classification

IBM
Hadoop

NETFLIX
Recommendation

Yandex
Search Rankings

Applications

- Image Detection
- Face Recognition
- Gesture Recognition
- Video Search & Analytics
- Speech Recognition & Translation
- Recommendation Engines
- Indexing & Search

Talks @ GTC

facebook

NYU

STANFORD UNIVERSITY

DARPA

DENSO

Carnegie Mellon University

MIT Massachusetts Institute of Technology

Berkeley UNIVERSITY OF CALIFORNIA

NVIDIA Invited to Participate in CI2015

NVIDIA

NCAR UCAR | CISL Computational & Information Systems Lab | Software | Data Portals | User Support | Resources | About Us

CI 2015

5th International Workshop on Climate Informatics
September 24-25, 2015

Hosted by the National Center for Atmospheric Research in Boulder, CO

CI2015 - Agenda

September 25, 2015, NCAR, MSR

8:30 – 9:30	**Invited Talk – Rich Caruana:** *Do Deep Nets Really Need To Be Deep?*
9:30 – 10:15	Panel Discussion: Deep Learning for Climate Science
10:45 – 11:45	**Invited Talk – Imme Ebert-Uphoff:** *Knowledge discovery in climate science*
1:00 – 2:00	**Invited Talk – Yolanda Gil:** *Intelligent Systems for Climate Research: When Will Deep Learners Meet Deep Knowledge?*

5th Intl Workshop on Climate Informatics, 24 – 25 Sep, 2015

NCAR, Boulder, USA

Growing Agenda on Deep Learning

Index: Scalable Rendering for Volume Visualization

- **Leverages GPU-clusters for large-scale (volume) data visualization and interactive visual computing**
- **Commercial software solution available and deployed for in-situ visualization of large-scale data**
- **Plugin for ParaView under development and available soon**
- http://www.nvidia-arc.com/products/nvidia-index.html

1.8 billion cells + 500 time steps

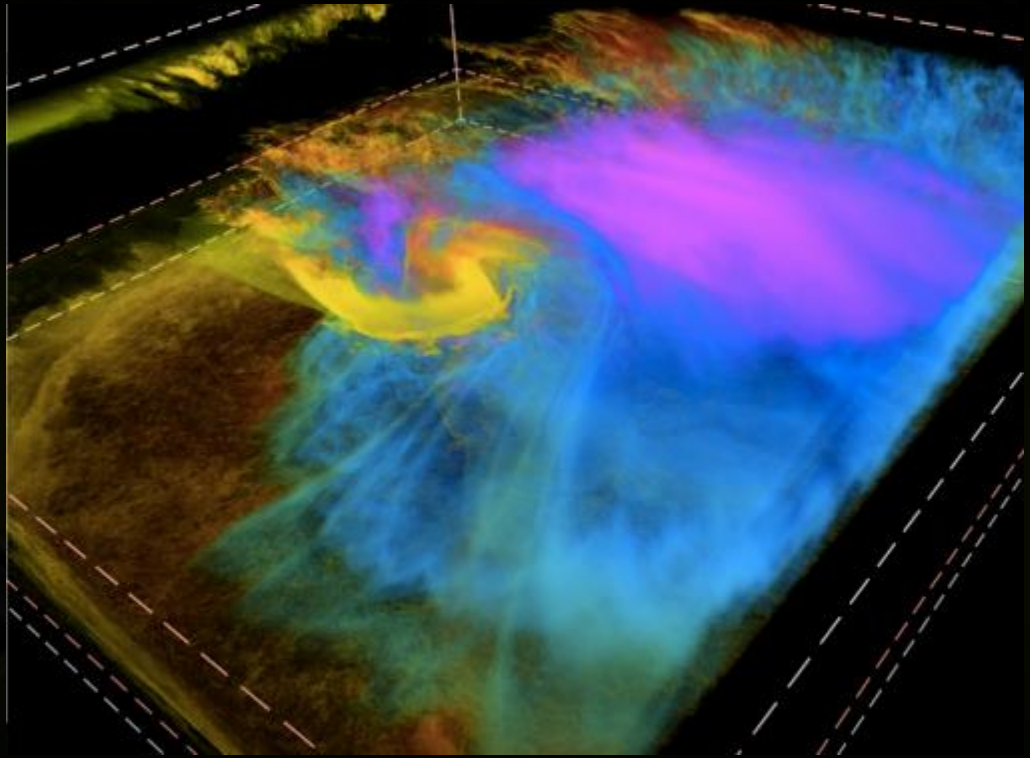

Dataset courtesy of Prof. Leigh Orf, UW-Madison and Rob Sisneros, NCSA

Agenda: Application Readiness for the Pre-Exascale Phase

- NVIDIA Update and HPC Trends
- **Developments in Pre-Exascale HPC**
- ESM Requirements and GPU Progress

IBM Power + NVIDIA GPU Accelerated HPC

NVIDIA

Next-Gen IBM Supercomputers and Enterprise Servers

Long term roadmap integration

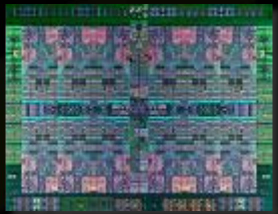

POWER
CPU

+

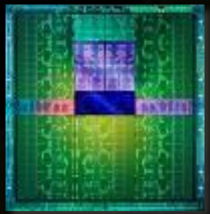

Tesla
GPU

OpenPOWER Foundation

Open ecosystem built on Power Architecture

NVIDIA

IBM®

Google

Mellanox TECHNOLOGIES

TYAN

& 30+ more…

First GPU-Accelerated POWER-Based Systems Available Since 2015

GPU and CPU Platforms and Configurations

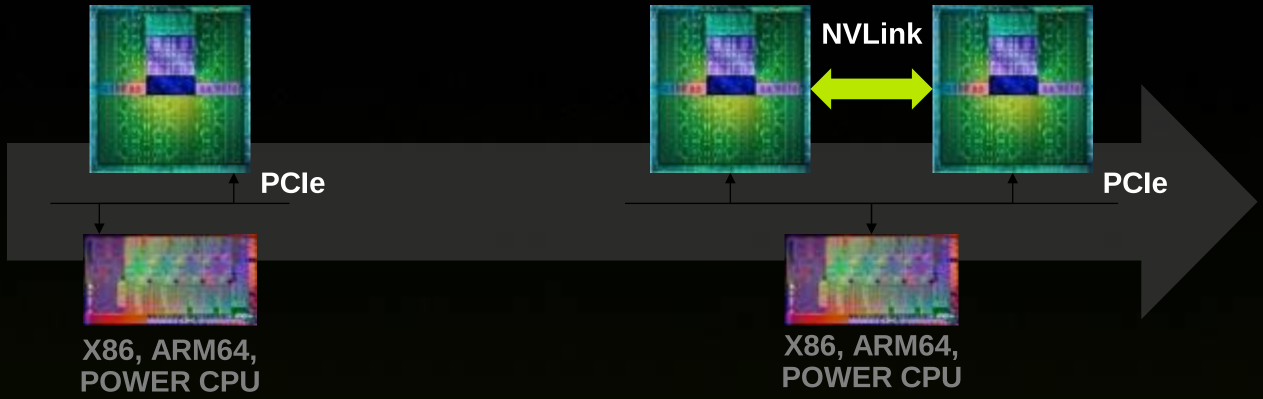

2014 – Kepler GPU

2016 – Pascal GPU

- More CPU Choices for GPU Acceleration – 2014
- NVLink Interconnect Alternative to PCIe – 2016

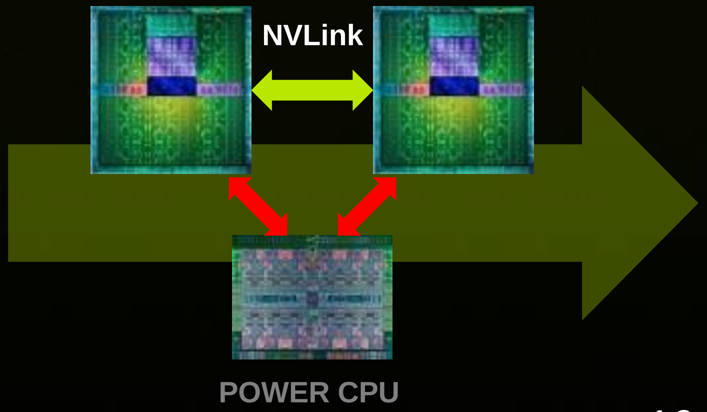

Features of Pascal GPU Architecture — 2016

NVLink
Interconnect at 80 GB/s
(Speed of CPU Memory)

Stacked Memory
4x Higher Bandwidth ~1 TB/s
3x Capacity, 4x More Efficient

Unified Memory
Lower Development Effort
(Available Today in CUDA6)

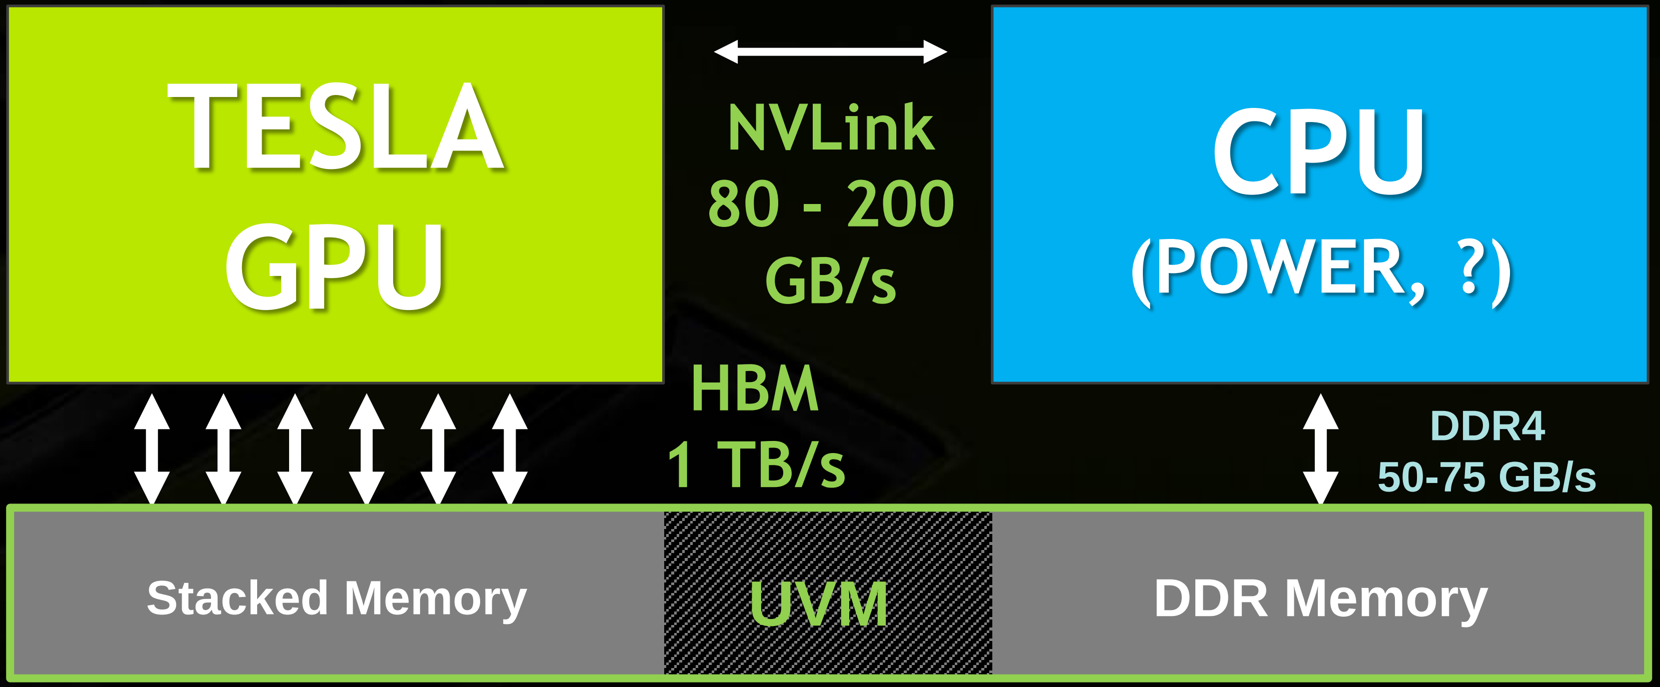

NVLink Interconnect Alternative to PCIe

- High speed interconnect for CPU-to-GPU and GPU-to-GPU
- Performance advantage over PCIe Gen-3 (5x to 12x faster)
- GPUs and CPUs share data structures at CPU memory speeds

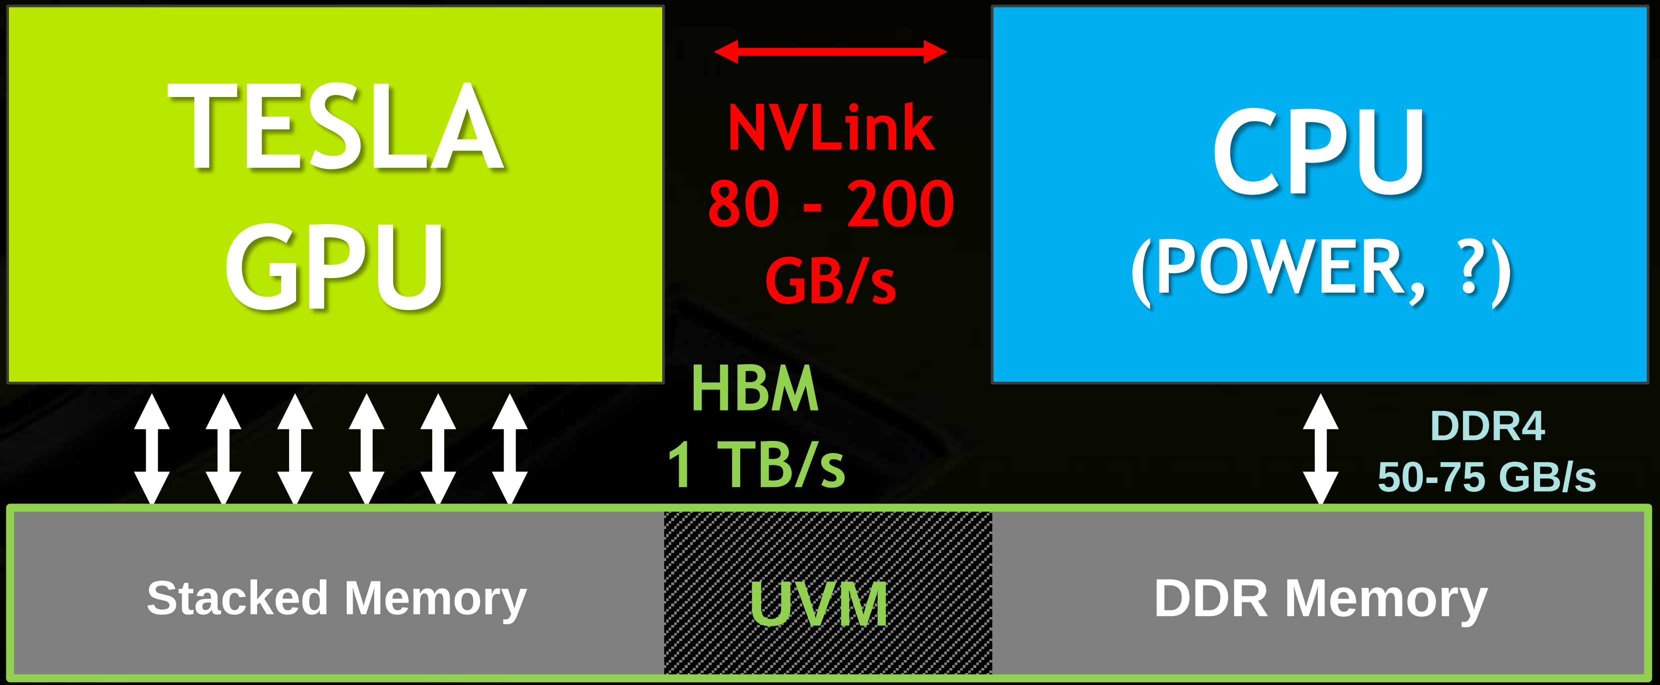

NVLink Interconnect Alternative to PCIe

NVIDIA

- GPU Density Increasing

CRAY THE SUPERCOMPUTER COMPANY

Cray CS-Storm: 8 x K80

Dell C4130: 4 x K80 DELL

hp

HP SL270: 8 x K80

- HPC Vendors & NVLink—2016
  - All will support GPU-to-GPU
  - Several to support CPU-to-GPU

Bull atos technologies | CRAY | DELL | hp | IBM | QCT Quanta CLOUD TECHNOLOGY | TYAN

US DOE Pre-Exascale Phase Deployment in 2017

- **US DOE CORAL Systems**
  - Summit (ORNL) and Sierra (LLNL)
  - Installation 2017 at ~150 PF each
  - Nodes of POWER 9 + Tesla Volta GPUs
  - NVLink Interconnect for CPUs + GPUs
- **ORNL Summit System**
  - Approximately 3,400 total nodes
  - Each node 40+ TF peak performance
  - About 1/5 of total #2 Titan nodes (18K+)
  - Same energy used as #2 Titan (27 PF)

SUMMIT
Oak Ridge National Laboratory

SIERRA
Lawrence Livermore National Laboratory

CORAL
COLLABORATION
OAK RIDGE • ARGONNE • LIVERMORE

CORAL Summit System
5-10x Faster than Titan
1/5th the Nodes,
Same Energy Use as Titan

Summit System Relative to HPC Top 500

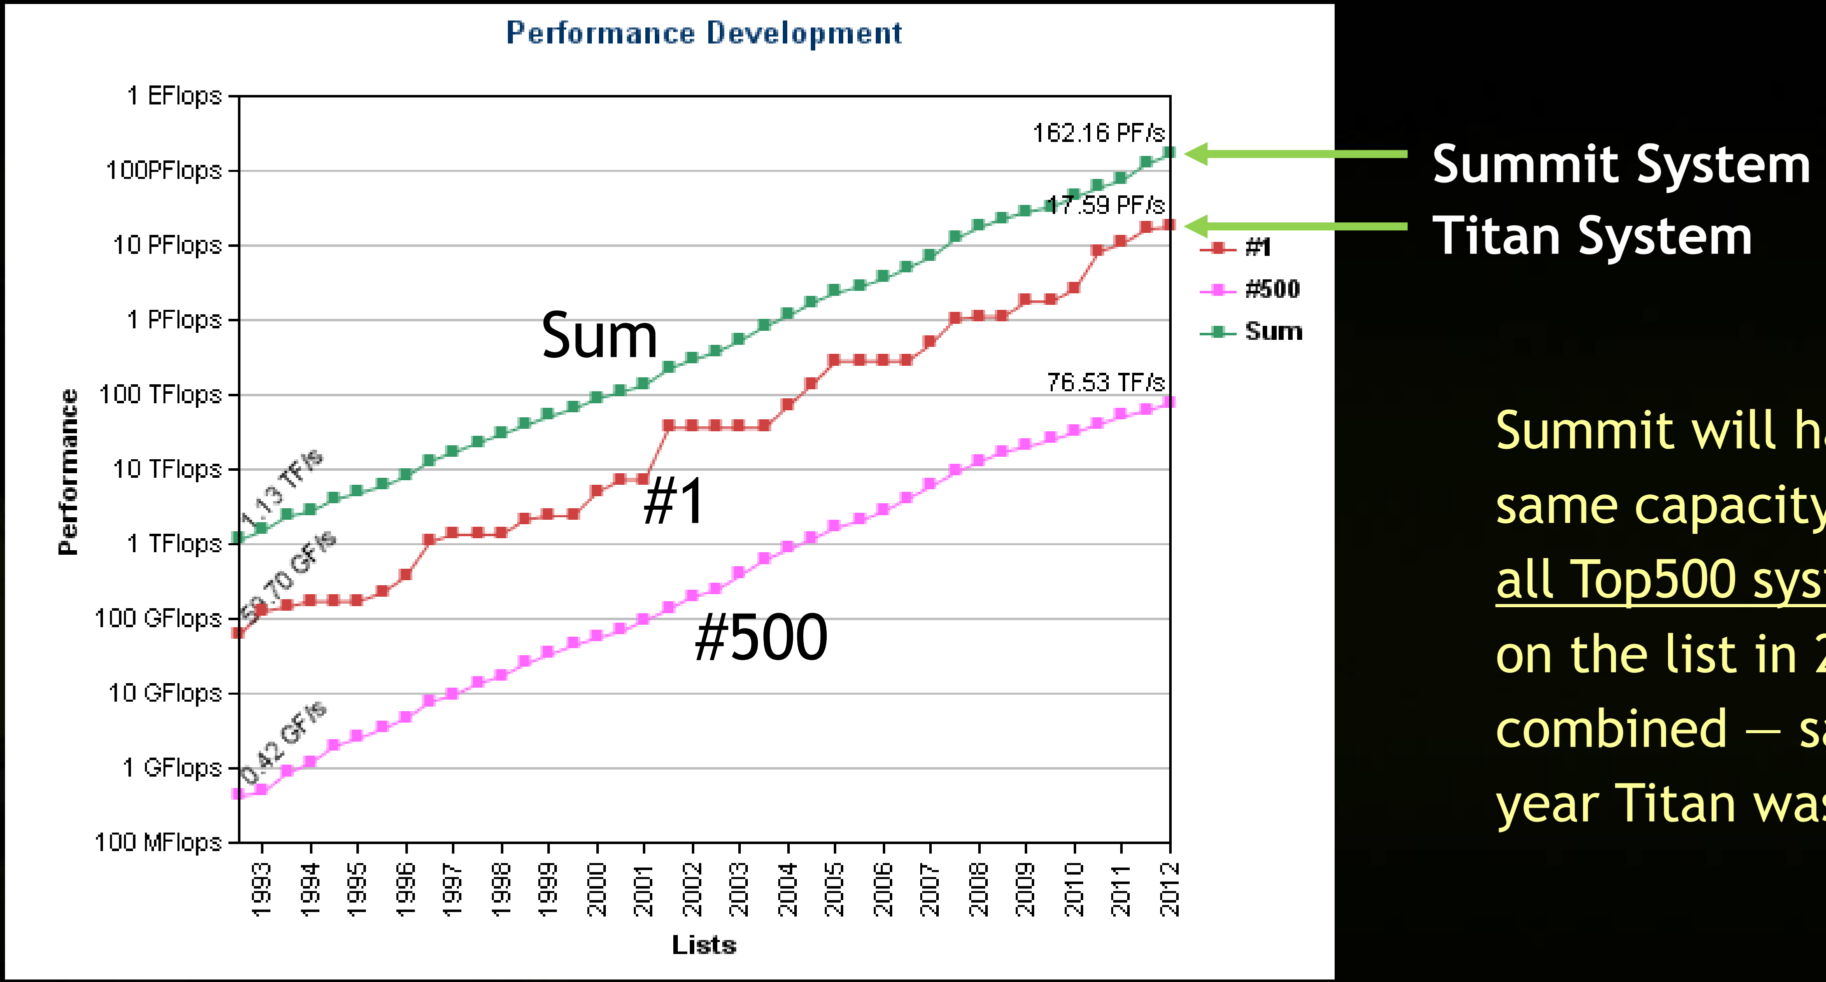

Summit will have same capacity as all Top500 systems on the list in 2012 combined – same year Titan was #1

Feature Comparison of Summit vs. Titan

Feature	2017 Summit	2012 Titan	
Application Performance	5-10x Titan	Baseline	~5 - 10x
Number of Nodes	~3,400	18,688	~1/5x
Node performance	> 40 TF	1.4 TF	~30x
Memory per Node	>512 GB (HBM + DDR4)	38GB (GDDR5+DDR3)	~15x
NVRAM per Node	800 GB	0	
Node Interconnect	NVLink (5-12x PCIe 3)	PCIe 2	~10 - 20x
System Interconnect (node injection bandwidth)	Dual Rail EDR-IB (23 GB/s)	Gemini (6.4 GB/s)	
Interconnect Topology	Non-blocking Fat Tree	3D Torus	
Processors	IBM POWER9 NVIDIA Volta™	AMD Opteron™ NVIDIA Kepler™	
File System	120 PB, 1 TB/s, GPFS™	32 PB, 1 TB/s, Lustre®	
Peak power consumption	10 MW	9 MW	~0x

US DOE Projections of Exascale by 2022

Present and Future Leadership Computers at OCLF – J. Wells, GTC 2015

Agenda: Application Readiness for the Pre-Exascale Phase

- NVIDIA Update and HPC Trends
- Developments in Pre-Exascale HPC
- **ESM Requirements and GPU Progress**

ESM Characteristics and GPU Considerations

- Flat execution profiles – no real ‘hot spots’
  - Programming models must enable optimization-everywhere for GPUs
- Large Fortran code base with a ‘community’ of developers
  - Usually restricts some portion of the code/model to directives approach
- Dynamical cores distributed parallel with frequent communication
  - Fast interconnects, GPU-aware MPI with direct GPU-to-GPU communication
- Development environment for multi-level parallelism
  - Libraries, tools, debuggers, compilers for OpenMP and OpenACC
- Data center infrastructure and system mgmt (vs. card ‘plug-in’)
  - Vendor collaboration/integration (Cray, IBM, SGI, Bull, HP, Lenovo, etc.)

Programming Strategies for GPU Acceleration

Applications		
GPU Libraries	OpenACC Directives	Programming Languages
Provides Fast "Drop-In" Acceleration	GPU-acceleration in Standard Language (Fortran, C, C++)	Maximum Flexibility with GPU Architecture and Software Features

Increasing Development Effort →

NOTE: Many application developments include a combination of these strategies

NVIDIA GPU Focus for Atmosphere Models

Global

Organization	Location	Model	GPU Approach
ORNL, SNL	US	CAM-SE	OpenACC (migration from CUDA-F)
NOAA ESRL	US	FIM/NIM	OpenACC, F2C-ACC
NASA GSFC	US	GEOS-5	OpenACC (migration from CUDA-F)
NCAR MMM	US	MPAS-A	OpenACC
NOAA GFDL	US	FV3	OpenACC
ECMWF	UK	IFS (Arpege)	Libs + OpenACC
STFC, MetOffice	UK	UM/GungHo	OpenACC
CSCS, MPI-M	CH, DE	ICON	DSL – dycore, OpenACC – physics
JAMSTEC, UT, RIKEN	JP	NICAM	OpenACC

Regional

(i) NCAR; (ii) SSEC	US	WRF-ARW	(i) OpenACC, (ii) CUDA
MCH, CSCS	CH	COSMO	DSL – dycore, OpenACC – physics
Bull	FR	HARMONIE	OpenACC
TiTech	JP	ASUCA	Hybrid-Fortran, OpenACC

NVIDIA GPU Focus for Atmosphere Models

Global

Organization	Location	Model	GPU Approach
ORNL, SNL	US	CAM-SE	OpenACC (migration from CUDA-F)
NOAA ESRL	US	FIM/NIM	OpenACC, F2C-ACC
NASA GSFC	US	GEOS-5	OpenACC (migration from CUDA-F)
NCAR MMM	US	MPAS-A	OpenACC
NOAA GFDL	US	FV3	OpenACC
ECMWF	UK	IFS (Arpege)	Libs + OpenACC
STFC, MetOffice	UK	UM/GungHo	OpenACC
CSCS, MPI-M	CH, DE	ICON	DSL – dycore, OpenACC – physics
JAMSTEC, UT, RIKEN	JP	NICAM	OpenACC

Regional

(i) NCAR; (ii) SSEC	US	WRF-ARW	(i) OpenACC, (ii) CUDA
MCH, CSCS	CH	COSMO	DSL – dycore, OpenACC – physics
Bull	FR	HARMONIE	OpenACC
TiTech	JP	ASUCA	Hybrid-Fortran, OpenACC

US DOE Accelerator-Based Climate Model ACME

- ACME: Accelerated Climate Model for Energy
  - First fully accelerated climate model (GPU and MIC)
  - Consolidation of DOE ESM projects from 7 into 1
  - DOE Labs: Argonne, LANL, LBL, LLNL, ORNL, PNNL, Sandia

- ACME a development branch of CESM from NCAR
  - Atmosphere component CAM-SE (NCAR)
  - Ocean component MPAS-O (LANL)
  - Others: CICE (LANL), CLM (NCAR), SAM (CSU)
  - Towards NH global atm 12 km, ocn 15 km, 80 years

- Co-design project using US DOE LCF systems

First Report – 11 Jul 2014

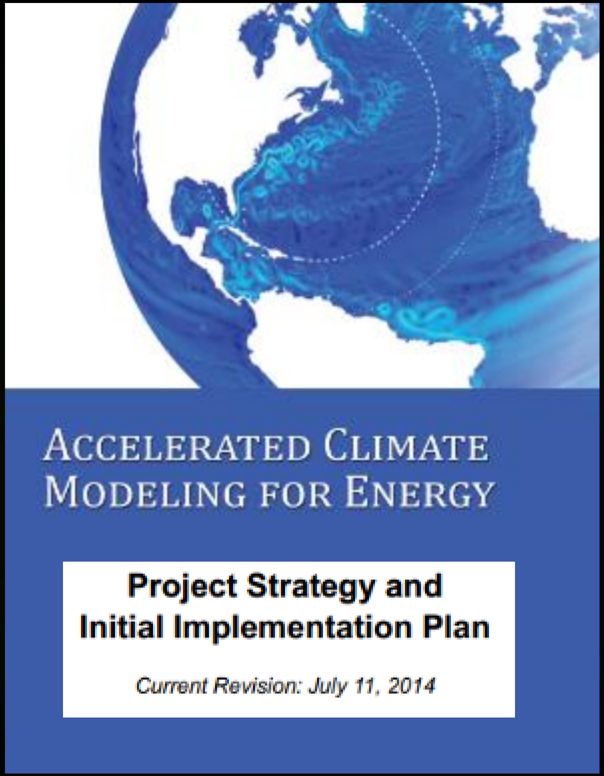

Source: http://climatemodeling.science.energy.gov/sites/default/files/publications/acme-project-strategy-plan_0.pdf

US DOE Accelerator-Based Climate Model ACME

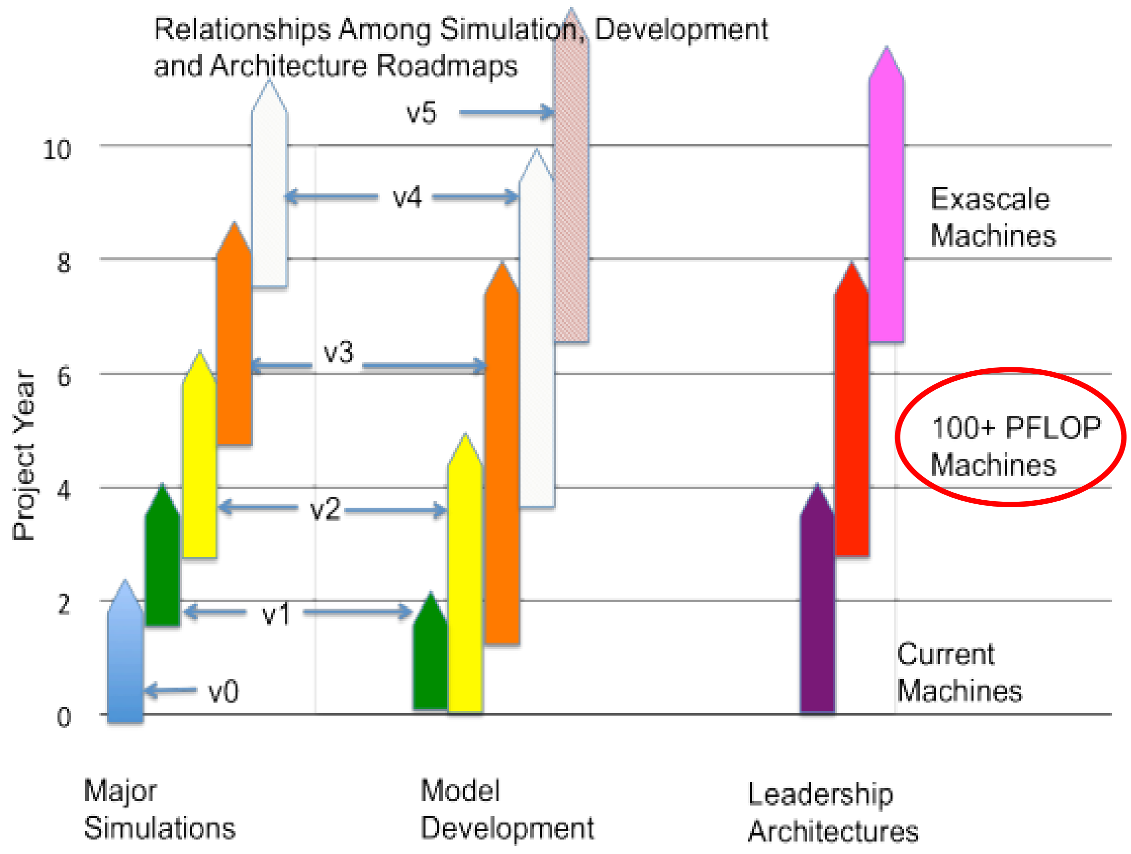

ACME Project Roadmap
(page 2 in report)

2022

- **TBD**

2017

- **CORAL: Summit, Sierra, Aurora**
- **Trinity: NERSC8 (Cori)**

2014

- **TITAN – OLCF** [AMD + GPU]
- **Mira – ALCF** [Blue GeneQ]

Source: http://climatemodeling.science.energy.gov/sites/default/files/publications/acme-project-strategy-plan_0.pdf

ICON ATM Climate Model and OpenACC

Current Status (May 2015)

PRACE 2IP extension has ended. So what is working?

- *Dynamical core:* 1 hr. R2B04 (4 GPUs) validates to 1.E-13.
- Implemented on GPU: *solve_nh*, *diffusion*, underlying routines, including *mo_sync* for communication
- Some advection schemes implemented (not in benchmark)
- Here: first publicly announced trunk NHDC benchmarks
- Actively porting 2-moment microphysics to OpenACC
- Radiation RRTMGP under discussion with Pincus, et al.

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

15

CSCS
Swiss National Supercomputing Centre

Tuesday, June 2, 15

PASC15 Conference
01 Jun 15, ETH Zurich

Using GPUs Productively for the ICON Climate Model
-by Dr. Will Sawyer, CSCS

ICON Climate Model and OpenACC on GPUs

NVIDIA

ICON Trunk NHDC on GPUs Results *(New!)*

Compare ICON Trunk Piz Dora (2x Haswell sockets) vs. Piz Daint nodes w/ K20x (both with Cray CCE)

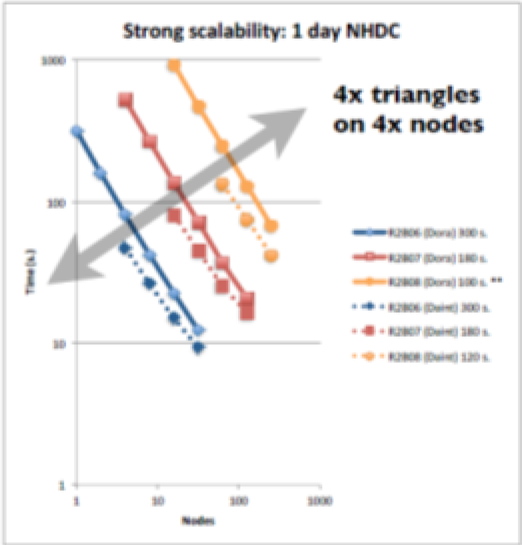

Lower is Better

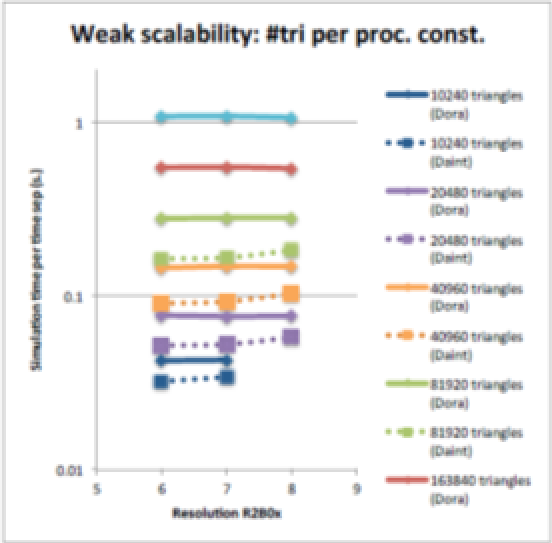

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

$$\frac{[\text{R2B07 Dora 64 (s)}]}{[\text{R2B07 Daint 64 (s)}]} = 1.45 \qquad \frac{[\text{BW Daint node (GB/s)}]}{[\text{BW Dora node (GB/s)}]} = \frac{165}{110} = 1.5$$

CSCS
Swiss National Supercomputing Centre

Tuesday, June 2, 15

PASC15 Conference
01 Jun 15, ETH Zurich

Using GPUs Productively for the ICON Climate Model
-by Dr. Will Sawyer, CSCS

- **Strong Scalability to 120 GPU nodes of Piz Daint**
- **Node comparisons of K20X vs. 2 x Haswell**
- **Non-hydrostatic dycore only, no nesting**

ECMWF Project for Exascale Weather Prediction

Expectations towards Exascale: Weather and Climate Prediction – P. Bauer, ECMWF, EESI-2015

ESCAPE*, Energy efficient SCalable Algorithms for weather Prediction at Exascale:

- Next generation IFS numerical building blocks and compute intensive algorithms
- Compute/energy efficiency diagnostics
- New approaches and implementation on novel architectures
- Testing in operational configurations

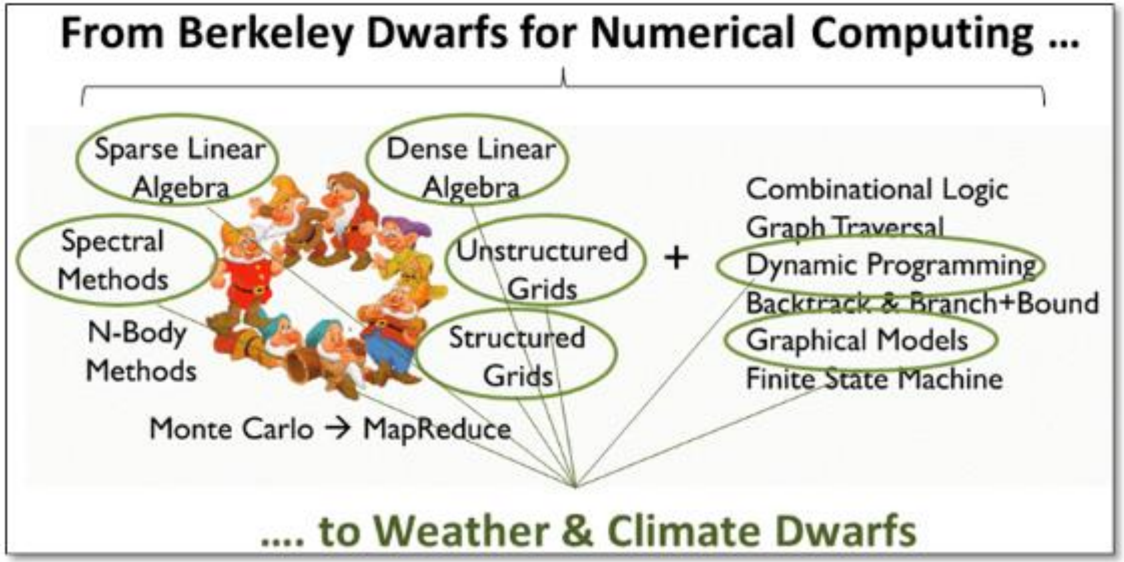

*To be funded by EC H2020 framework, Future and Emerging Technologies – High-Performance Computing; Partners: **ECMWF**, Météo-France, RMI, DMI, Meteo Swiss, DWD, Loughborough U, PSNC, Bull, NVIDIA, Optalysys

EASCAPE HPC Goals:

- Standardized, highly optimized kernels on specialized hardware
- Overlap of communication and computation
- Compilers/standards supporting portability

Project Approach:

- Co-design between domain scientists, computer scientists, and HPC vendors

ECMWF Initial IFS Investigation on Titan GPUs

Challenges of Getting ECMWF's Weather Forecast Model (IFS) to the Exascale
– G. Mozdzynski, ECMWF 16th HPC Workshop

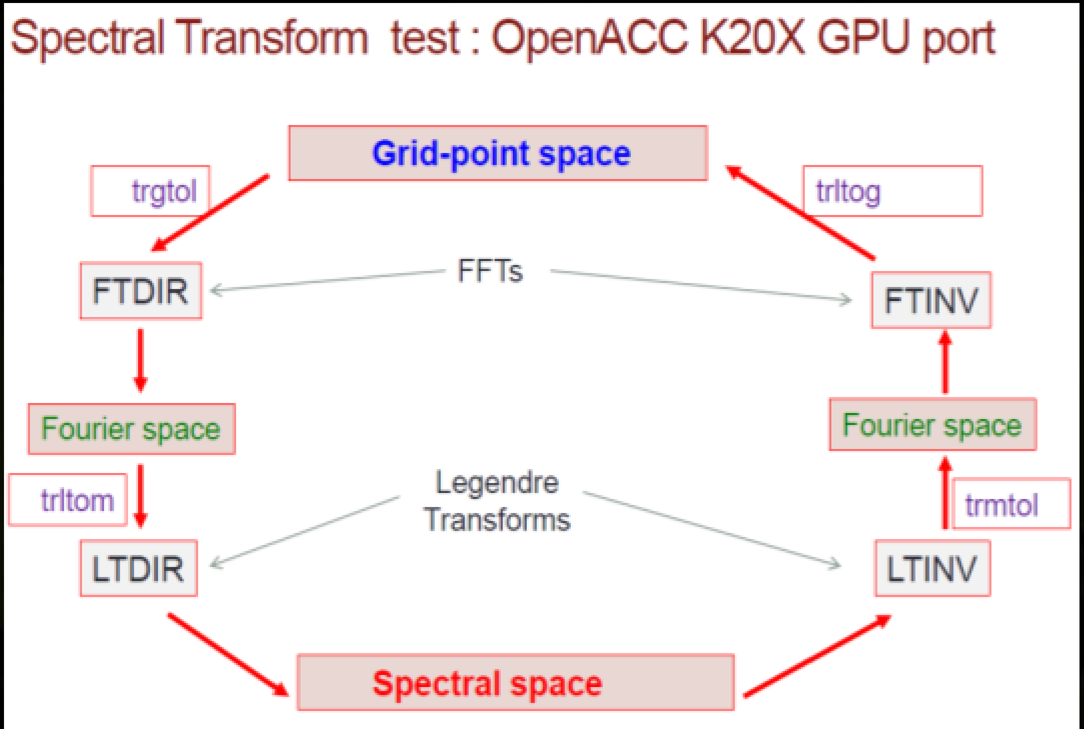

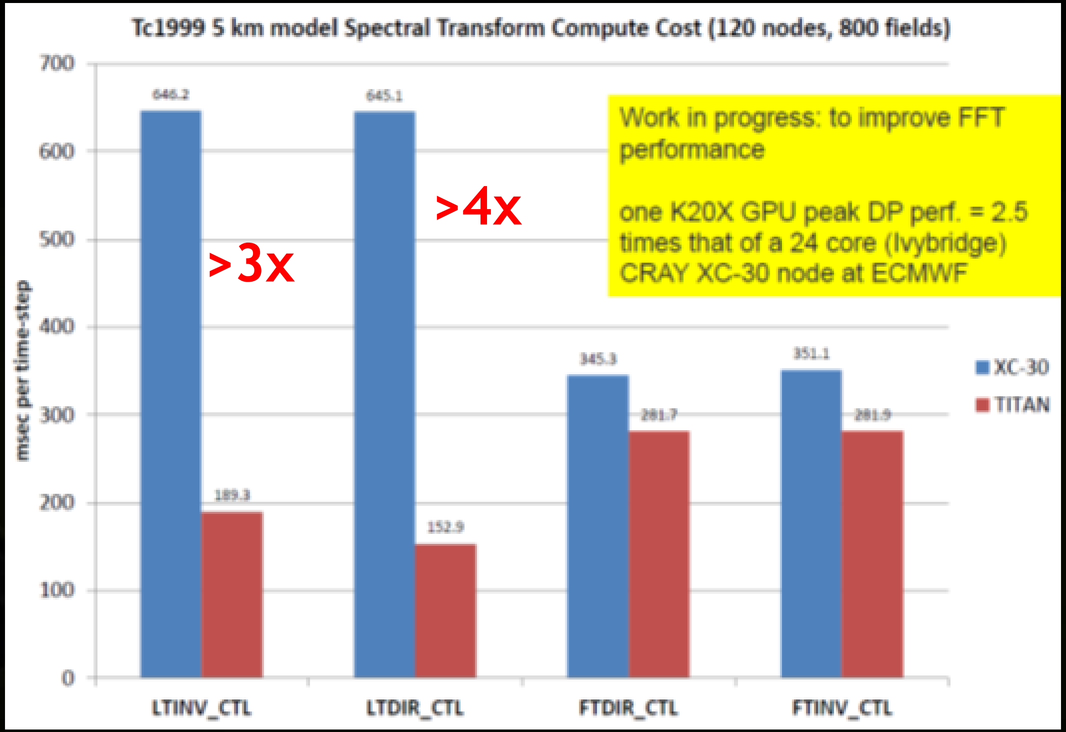

Legendre Transform Kernels Observe >4x for Titan K20X vs. XC-30 24-core Ivy Bridge

NOAA NGGPS: NH Model Dycore Candidates (5)

Model	Organization	Numeric Method	Grid
NIM	NOAA/ESRL	Finite Volume	Icosahedral
MPAS	NCAR/LANL	Finite Volume	Icosahedral/Unstructured
NEPTUNE	Navy/NRL	Spectral Element	Cubed-Sphere with AMR
HIRAM/FV3	NOAA/GFDL	Finite Volume	Cubed-Sphere, nested
NMM-UJ	NOAA/EMC	Finite difference	Cubed-Sphere
GFS-NH	NOAA/EMC	Semi-Lagrangian/Spectral	Reduced Gaussian
IFS (RAPS13)	ECMWF	Semi-Lagrangian/Spectral	Reduced Gaussian

From: Next Generation HPC and Forecast Model Application Readiness at NCEP
-by John Michalakes, NOAA NCEP; AMS, Phoenix, AZ, Jan 2015

Dycores FV3 & MPAS selected for Level-II evaluation – NVIDIA, NOAA investigating potential for GPUs

NOAA evaluations under HPC program SENA – Software Engineering for Novel Architectures (2015)

GFDL Participation in GPU Hackathon Oct 2015

- **Mentors and content are contributed by recognized HPC leaders:**
  - **USA:** Oak Ridge National Laboratory, NCSA, Cray, NVIDIA, PGI
  - **Europe:** CSCS, TU Dresden
  - **Japan:** Tokyo Institute of Technology

- **NOAA GFDL team led by Dr. Rusty Benson with NVIDIA and PGI members**

COSMO Regional Model on GPUs

http://www.cosmo-model.org/

- **COSMO model fully implemented on GPUs by MeteoSwiss**
  - COSMO consortium approved GPU code for standard distribution

- **MeteoSwiss to deploy first-ever operational NWP on GPUs**
  - Successful daily test runs on GPUs since Dec 2014, operational in 2016

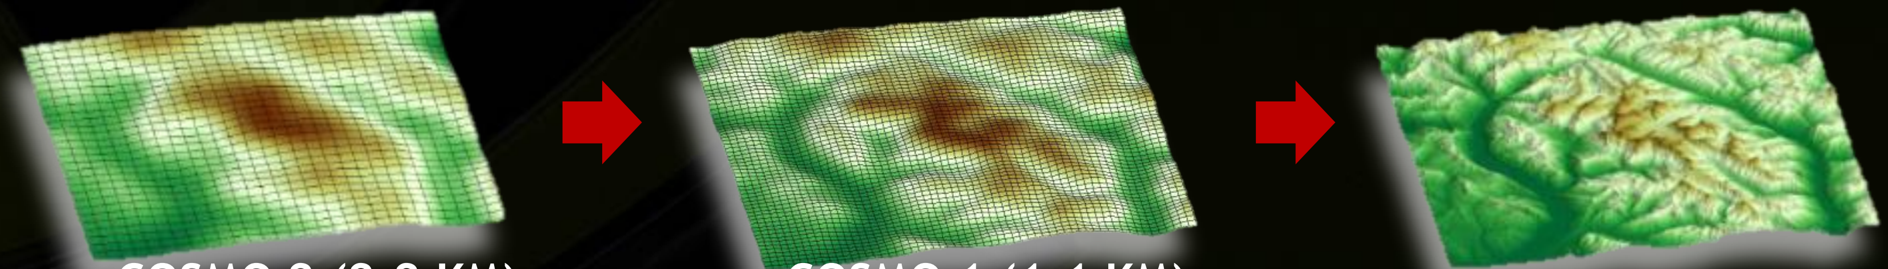

COSMO 2 (2.2 KM) → COSMO 1 (1.1 KM) → Reality

Source:

CSCS Swiss National Supercomputing Centre

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

"Increased Resolution = Quality Forecast"

WRF Progress Update

http://www.wrf-model.org/

- **Independent development projects of (I) CUDA and (II) OpenACC**

I. CUDA through funded collaboration with SSEC www.ssec.wisc.edu

- SSEC lead Dr. Bormin Huang, NVIDIA CUDA Fellow: research.nvidia.com/users/bormin-huang
- Objective: Hybrid WRF benchmark capability starting 2H 2015

II. OpenACC directives through collaboration with NCAR MMM

- Objective: NCAR GPU interest in Fortran version to host on distribution site

Project Status (II): NVIDIA-NCAR OpenACC WRF

- **Objective to provide benchmark capability of hybrid CPU+GPU WRF 3.6.1**
- **NVIDIA ported 30+ routines to OpenACC for both dynamics and physics:**
  - List provided on next page
- **Project to integrate OpenACC modules with full model WRF 3.6.1 release**
  - Open source project with plans for NCAR hosting on WRF trunk distribution site
  - Objective to minimize code refactoring for scientists’ readability and acceptance

NVIDIA - NCAR Project Success

- **Several performance sensitive routines of dynamics and physics are completed**
  - This includes advection routines for dynamics; expensive microphysics schemes
- **Hybrid WRF performance using single socket CPU + GPU faster than CPU-only**
  - Incremental improvements as more OpenACC code is developed for execution on GPU

WRF Modules Available in OpenACC

- **Project to implement OpenACC routines into full model WRF 3.6.1**
  - Several dynamics routines including all of advection
  - Several physics schemes:
    - Microphysics – Kessler, Morrison, Thompson
    - Radiation – RRTM
    - Planetary boundary layer – YSU, GWDO
    - Cumulus – KFETA
    - Surface physics – Noah

```
dyn_em/module_advect_em.OpenACC.F
dyn_em/module_bc_em.OpenACC.F
dyn_em/module_big_step_utilities_em.OpenACC.F
dyn_em/module_diffusion_em.OpenACC.F
dyn_em/module_em.OpenACC.F
dyn_em/module_first_rk_step_part1.OpenACC.F
dyn_em/module_first_rk_step_part2.OpenACC.F
dyn_em/module_small_step_em.OpenACC.F
dyn_em/module_stoch.OpenACC.F
dyn_em/solve_em.OpenACC.F
dyn_em/start_em.OpenACC.F
frame/module_dm.OpenACC.F
frame/module_domain_extra.OpenACC.F
frame/module_domain.OpenACC.F
frame/module_domain_type.OpenACC.F
phys/module_bl_gwdo.OpenACC.F
phys/module_bl_ysu.OpenACC.F
phys/module_cu_kfeta.OpenACC.F
phys/module_microphysics_driver.OpenACC.F
phys/module_microphysics_zero_out.OpenACC.F
phys/module_mp_kessler.OpenACC.F
phys/module_mp_morr_two_moment.OpenACC.F
phys/module_mp_thompson.OpenACC.F
phys/module_pbl_driver.OpenACC.F
phys/module_physics_addtendc.OpenACC.F
phys/module_physics_init.OpenACC.F
phys/module_ra_rrtm.OpenACC.F
phys/module_ra_sw.OpenACC.F
phys/module_sf_noahlsm.OpenACC.F
phys/module_sf_sfclayrev.OpenACC.F
share/module_bc.OpenACC.F
share/wrf_bdyin.OpenACC.F
```

Project Status (I): NVIDIA-SSEC CUDA WRF

- Objective to provide Q3 benchmark capability of hybrid CPU+GPU WRF 3.6.1
- SSEC ported 20+ modules to CUDA across 5 physics types in WRF:
  - Cloud MP = 11; Radiation = 4; Surface Layer = 3; Planetary BL = 1; Cumulus = 1
- Project to integrate CUDA modules into latest full model WRF 3.6.1
  - SSEC development began 2010, based on WRF 3.3/3.4 vs. latest release 3.6.1
  - Start with existing modules most common across target customer's input
  - Objective: speedups measured on individual modules within full model run

NVIDIA – SSEC Project Success

- Target modules achieve published GPU performance within full WRF model run
  - Final clean-up and performance studies underway, real benchmarks in 2H 2015
- TQI/SSEC to fund/port all modules to CUDA towards a full WRF model on GPUs
  - Focus on additional physics modules, and ongoing WRF-ARW dycore development

SSEC Speedups of WRF Physics Modules

* Modules for initial NVIDIA-funded integration project

WRF Module	GPU Speedup (w-wo I/O)	Technical Paper Publication
Kessler MP	**70x / 816x**	J. Comp. & GeoSci., 52, 292-299, 2012
Purdue-Lin MP	**156x / 692x**	SPIE: doi:10.1117/12.901825
WSM 3-class MP	**150x / 331x**	
WSM 5-class MP *	**202x / 350x**	JSTARS, 5, 1256-1265, 2012
Eta MP	**37x / 272x**	SPIE: doi:10.1117/12.976908
WSM 6-class MP *	**165x / 216x**	Submitted to J. Comp. & GeoSci.
Goddard GCE MP	**348x / 361x**	Accepted for publication in JSTARS
Thompson MP *	**76x / 153x**	
SBU 5-class MP	**213x / 896x**	JSTARS, 5, 625-633, 2012
WDM 5-class MP	**147x / 206x**	
WDM 6-class MP	**150x / 206x**	J. Atmo. Ocean. Tech., 30, 2896, 2013
RRTMG LW *	**123x / 127x**	JSTARS, 7, 3660-3667, 2014
RRTMG SW *	**202x / 207x**	Submitted to J. Atmos. Ocean. Tech.
Goddard SW	**92x / 134x**	JSTARS, 5, 555-562, 2012
Dudhia SW *	**19x / 409x**	
MYNN SL	**6x / 113x**	
TEMF SL	**5x / 214x**	
Thermal Diffusion LS	**10x / 311x**	Submitted to JSATRS
YSU PBL *	**34x / 193x**	Submitted to GMD

Hybrid WRF Customer Benchmark Capability Starting in 2H 2015

Hardware and Benchmark Case

CPU: Xeon Core-i7 3930K,1 core use;

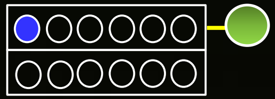

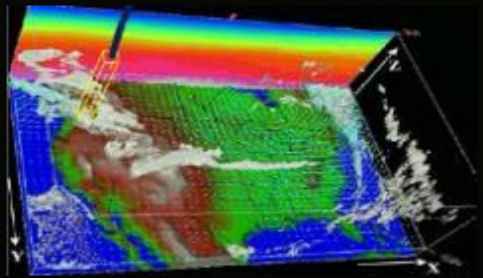

Benchmark: CONUS 12 km for 24 Oct 24; 433 x 308, 35 levels

TempoQuest USA Start-up to Commercialize WRF

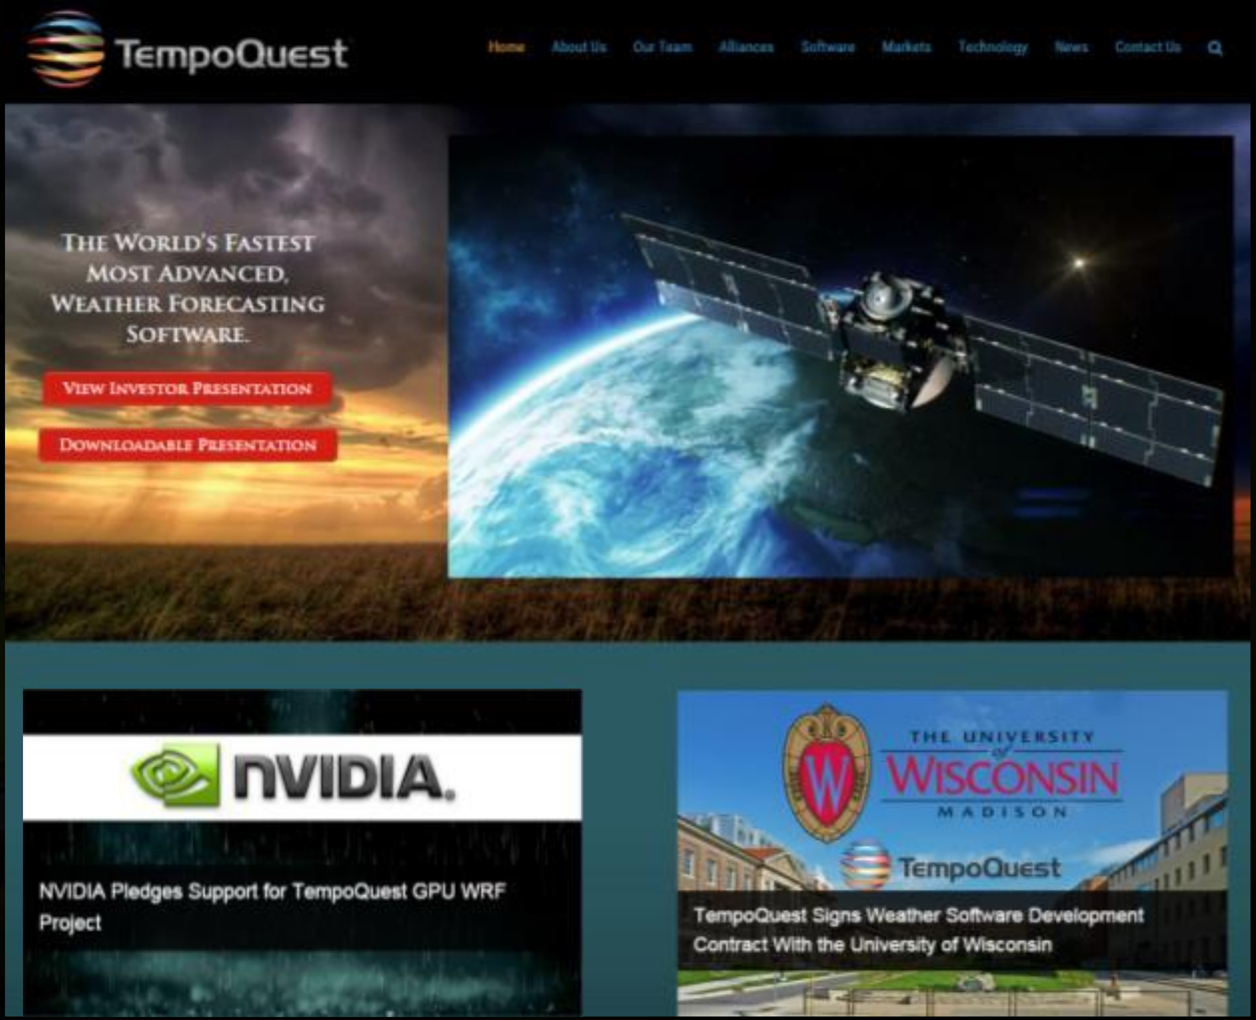

TQI funds SSEC to complete all WRF modules to CUDA

Focus on ARW dycore and additional physics

NVIDIA member of CUDA WRF project

Summary: Application Readiness for the Pre-Exascale Phase

- **NVIDIA observes strong ESM community interest in GPU acceleration**
  - Appeal of new technologies: Pascal, NVLink, more CPU platform choices
  - NVIDIA applications engineering collaboration in several model projects

- **NVIDIA HPC technology behind the next-gen pre-exascale systems**
  - New GPU technologies combined with OpenPOWER for DOE CORAL

- **GPU progress for several models – we examined a few of these**
  - Climate models DOE ACME and MPI-M/CSCS ICON ATM
  - NWP developments with ECMWF and ESCAPE, NOAA NGGPS, COSMO, WRF

Thank You

Stan Posey; sposey@nvidia.com; NVIDIA, Santa Clara, CA, USA