NVIDIA HPC Directions for Earth System Modeling

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

Agenda: NVIDIA HPC Directions for ESM

- NVIDIA HPC and GPU Update
- ES Model Progress on GPUs
- GPU Technology Roadmap

NVIDIA - Core Technologies and Products

Company Revenue of ~$5B USD; ~8,500 Employees; HPC Growing > 35% CAGR

GPU

GeForce®
Quadro®, Tesla®

Mobile

Tegra®

Cloud

GRID

GPUs Mainstream Across Diverse HPC Markets

NVIDIA

FY14 Segments

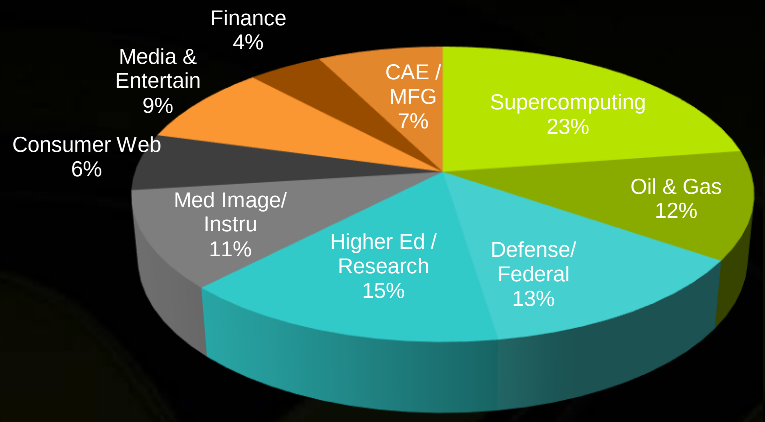

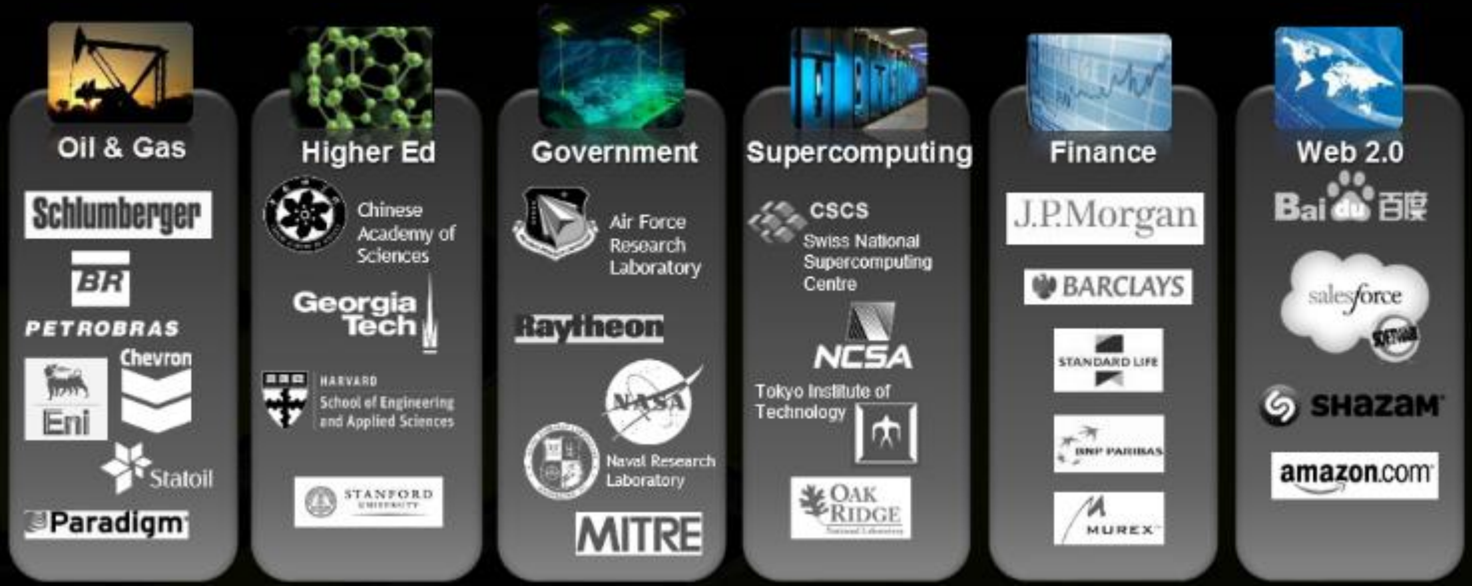

World's Top 3 Servers are GPU-Accelerated

IBM Power + NVIDIA GPU Accelerated HPC

Next-Gen IBM Supercomputers and Enterprise Servers

Long term roadmap integration

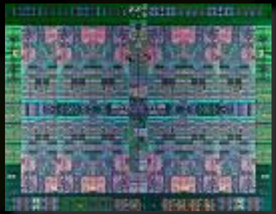

POWER
CPU

+

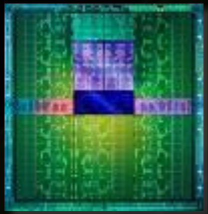

Tesla
GPU

OpenPOWER Foundation

Open ecosystem built on Power Architecture

& 30+ more…

First GPU-Accelerated POWER-Based Systems Available in Oct 2014

Cray + NVIDIA GPU Accelerated HPC

- **Cray and NVIDIA Collaboration on Large Systems**
  - **TITAN** – Oak Ridge National Laboratory, **#2 Top 500**
  - **Gaea** – NOAA (managed on-site at Oak Ridge)
  - **Piz Daint** – Swiss Supercomputing Center, **#6 Top 500**
  - **Blue Waters** – National Center for Scientific Applications

CRAY
THE SUPERCOMPUTER COMPANY

News Release

Cray Launches New High Density Cluster Packed With NVIDIA GPU Accelerators

SEATTLE, WA -- (Marketwired) -- 08/26/14 -- Global supercomputer leader Cray Inc. (NASDAQ: CRAY) today announced the launch of the Cray CS-Storm -- a high-density accelerator compute system based on the Cray® CS300™ cluster supercomputer. Featuring up to eight NVIDIA® Tesla® GPU accelerators and a peak performance of more than 11 teraflops per node, the Cray CS-Storm system is one of the most powerful single-node cluster architectures available today.

16 Aug 2014

Cray Launches High Density CS-Storm with NVIDIA K40 GPUs

- Up to 8 x K40 per node
- Capable of 1 PF per 4 Racks

Source: http://investors.cray.com/phoenix.zhtml?c=98390&p=irol-newsArticle&ID=1960850

GPU Motivation (I): Performance Trends

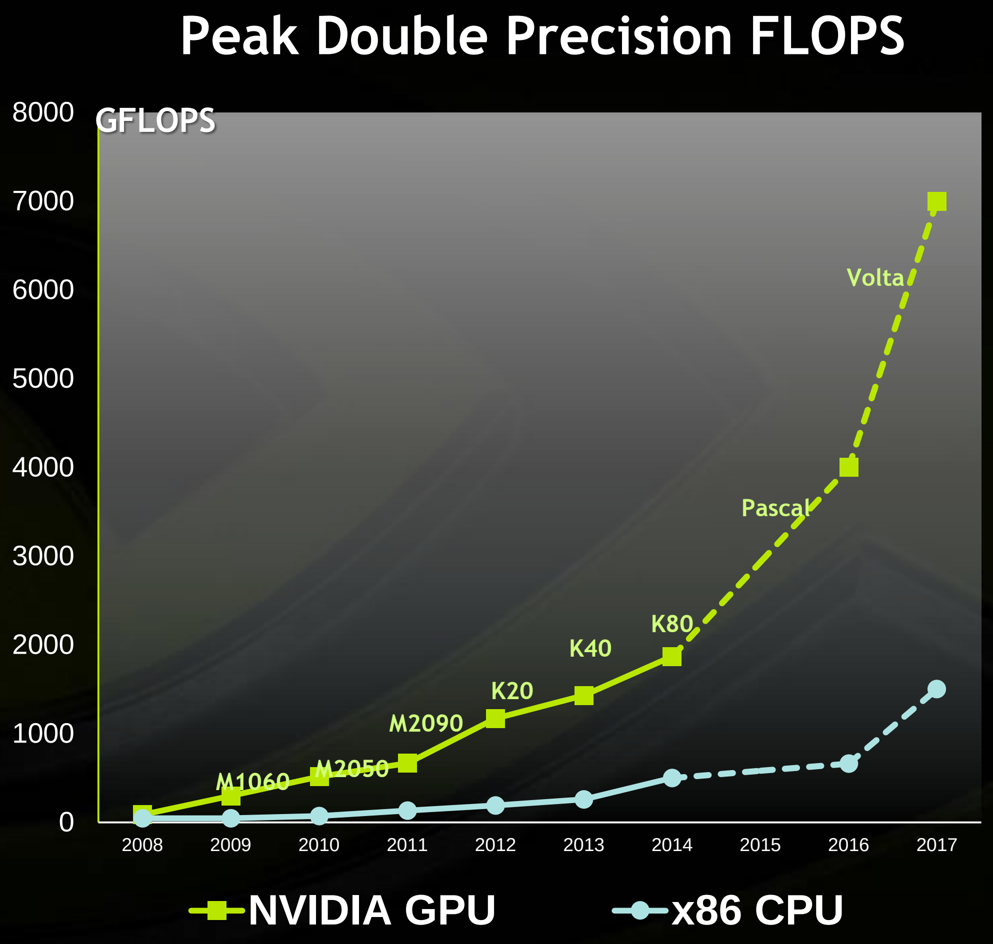

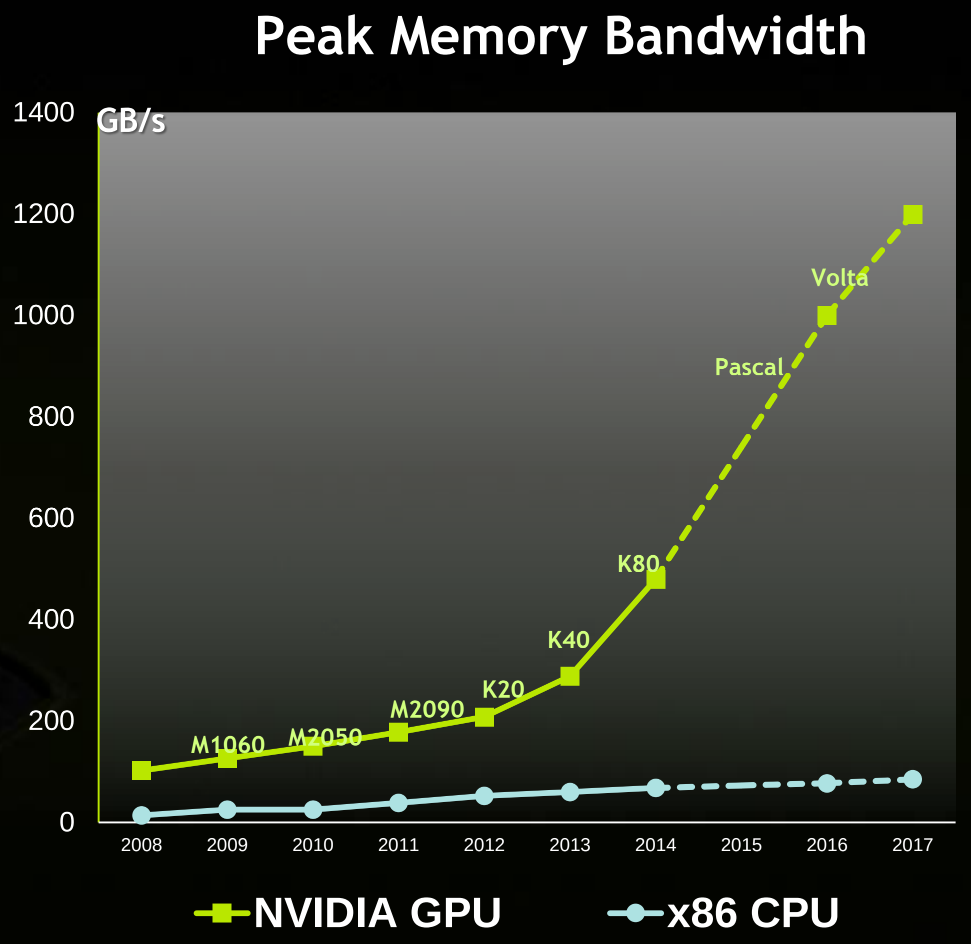

GPU Motivation (II): Energy Efficient HPC

NVIDIA

Top500 Rank	TFLOPS/s	Site
1	33,862.7	National Super Computer Centre Guangzhou
2	17,590.0	Oak Ridge National Lab **#1 USA**
3	17,173.2	DOE, United States
4	10,510.0	RIKEN Advanced Institute for Computational Science
5	8,586.6	Argonne National Lab
6	6,271.0	Swiss National Supercomputing Centre (CSCS) **#1 Europe**
7	5,168.1	University of Texas
8	5,008.9	Forschungszentrum Juelich
9	4,293.3	DOE, United States
10	3,143.5	Government

TOP500 SUPERCOMPUTER SITES

Green500 Rank	MFLOPS/W	Site
1	4,389.82	GSIC Center, Tokyo Tech **KFC**
2	3,631.70	Cambridge University
3	3,517.84	University of Tsukuba
4	3,459.46	SURFsara
5	3,185.91	Swiss National Supercomputing (CSCS)
6	3,131.06	ROMEO HPC Center
7	3,019.72	CSIRO
8	2,951.95	GSIC Center, Tokyo Tech **2.5**
9	2,813.14	Eni
10	2,629.10	(Financial Institution)
16	2,495.12	Mississippi State (top non-NVIDIA) **Intel Phi**
59	1,226.60	ICHEC (top X86 cluster)

THE GREEN 500

GPU Motivation (III): Model Trends in ESM

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

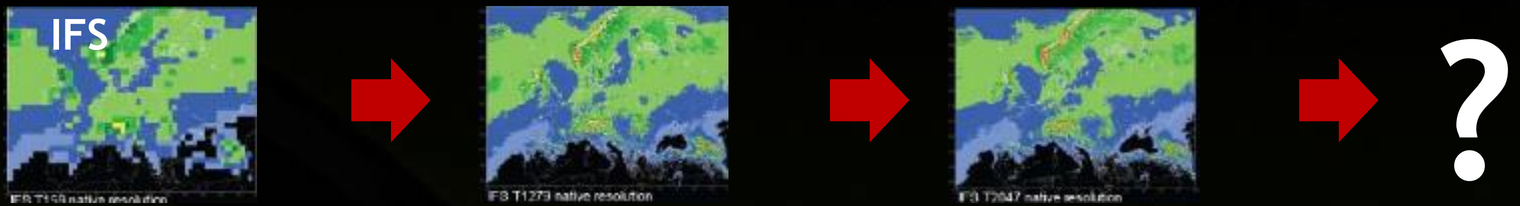

128 km | 16 km | 10 km | 3 km

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

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

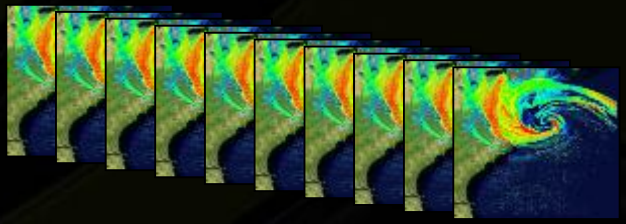

Number of Jobs > 10x

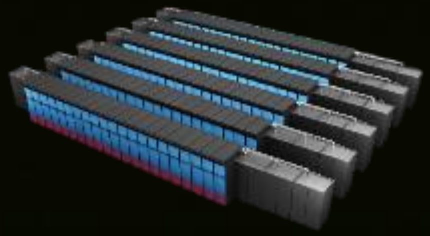

- **Fewer model approximations, more features (physics, chemistry, etc.)**

Accelerator technology identified as a cost-effective and practical approach to future computational challenges

Hardware Trends: NWP/Climate HPC Centers

NVIDIA

Motivation for x86 Migration Includes Preparation for Future Accelerator Deployment

Operational NWP

Organization	Location	Models	Previous/Current Operational HPC	Current/Next Operational HPC
ECMWF	Reading, UK	IFS	IBM Power	Cray XC30 – x86
Met Office	Exeter, UK	UM	IBM Power	Cray XC30 – x86
DWD	Offenbach, DE	GME, COSMO, ICON	NEC SX-9	Cray XC30 - x86
MF	Toulouse, FR	ALADIN, AROME	NEC SX-9	Bull - x86
NOAA/NCEP	*Various*, US	GFS, WRF, FIM, NIM	IBM Power	IBM iDataPlex - x86

Research

NCAR	Boulder, US	CESM, WRF, MPAS	IBM Power	IBM iDataPlex - x86
DKRZ/MPI-M	Hamburg, DE	MPI-ESM	IBM Power	Bull - x86

NVIDIA GPU Technology and ESM Strategy

Technology

- Development of GPUs as a co-processing accelerator for CPUs

Strategy

- Investments and alliances to develop Fortran-based programming environment
- Collaborations that provides applications engineering support in 16 model projects
- GPU integration and support on large systems with vendors (Cray, IBM, Bull, etc.)

New CPU Platforms Available During 2014

Tesla K40 (Kepler)

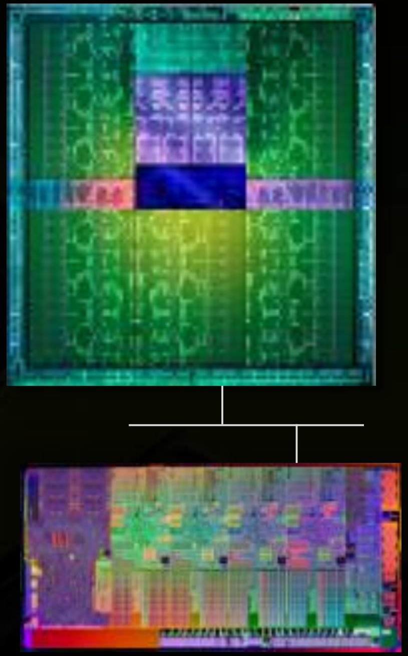

PCIe - 16 GB/s

X86 | ARM64 | Power8

Tesla GPU Progress Since 15th HPC Workshop

	ECMWF 15th Workshop **M2075**	**K20X**	**K40**	**15th / 16th Progress**
Peak SP Peak SGEMM	1.03 TF	3.93 TF 2.95 TF	4.29 TF 3.22 TF	**~ 4x**
Peak DP Peak DGEMM	.515 TF	1.31 TF 1.22 TF	1.43 TF 1.33 TF	**~3x**
Memory size	6 GB	6 GB	**12 GB**	**2x**
Mem BW (ECC off)	150 GB/s	250 GB/s	**288 GB/s**	**~2x**
Memory Clock		2.6 GHz	3.0 GHz	
PCIe Gen	Gen 2	Gen 2	Gen 3	**~2x**
# of Cores	448	2688	2880	**~5x**
Core Clock		732 MHz	Base: 745 MHz Boost: 875 Mhz	
Total Board Power	**235W**	**235W**	**235W**	**Same**

Agenda: NVIDIA HPC Directions for ESM

- NVIDIA HPC and GPU Update
- ES Model Progress on GPUs
- GPU Technology Roadmap

Tesla K40 Results from Multi-Core 4 Workshop

NVIDIA

NCAR UCAR CISL Computational & Information Systems Lab

2014 Heterogeneous Multi-Core 4 Workshop

September 17-18, 2014 at the National Center for Atmospheric Research in Boulder, Colorado
https://www2.cisl.ucar.edu/heterogeneous-multi-core-4-workshop/2014

1. **Modernizing Legacy Codes**
   - Youngsung Kim, NCAR CISL; **Source:** https://www2.cisl.ucar.edu/sites/default/files/youngsung_1.pdf

2. **Directive-Based Parallelization of the NIM**
   - Mark Govett, NOAA ESRL; **Source:** https://www2.cisl.ucar.edu/sites/default/files/govett_3.pdf

3. **Optimizing Weather Model Radiative Transfer Physics for the Many Integrated Core and GPGPU Architectures**
   - John Michalakes, NOAA NCEP; **Source:** https://www2.cisl.ucar.edu/sites/default/files/michalakes_1.pdf

4. **CUDA WRF Development Project at SSEC (presented at NVIDIA roadmap review)**
   - Jarno Mielikainen, SSEC UW-Madison; **Source:** contact Stan Posey, sposey@nvidia.com for full presentation

1. Tesla K40 Results: NCAR-CISL DG Kernels

We've learned some about

- how to get performance from new processors

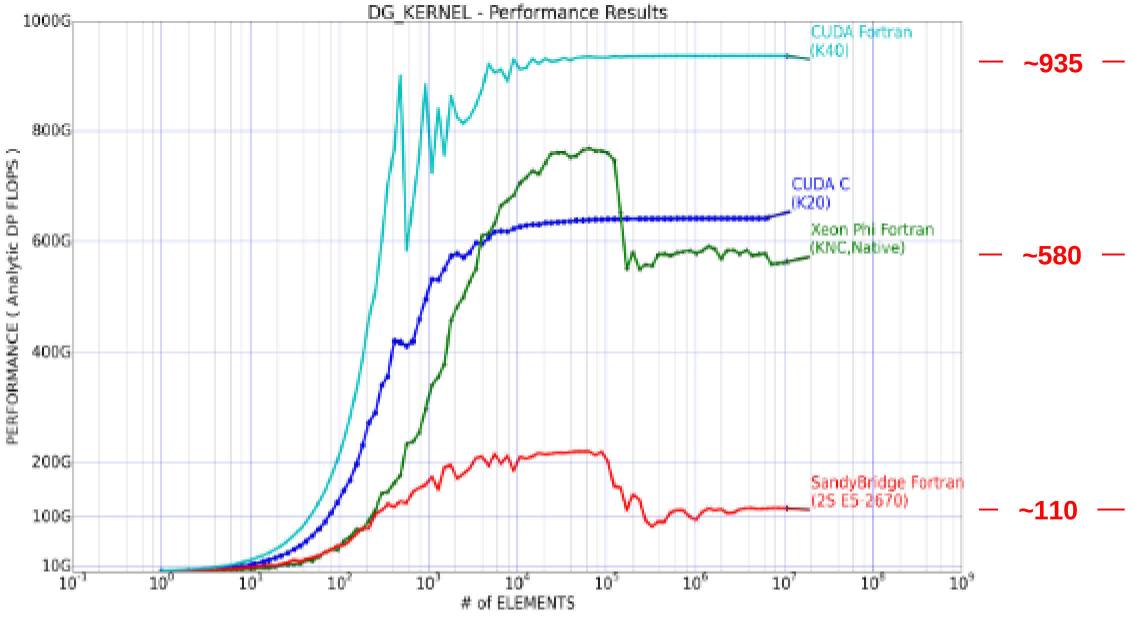

SUMMARY

K40 Speed-ups

Sandy Bridge	8.5x
Xeon Phi (Native)	1.6x

NOTES:

- Results for Number of Elements = 10^7
- K40 results from optimized CUDA Fortran

2. Tesla K40 Results: NOAA-ESRL and NIM

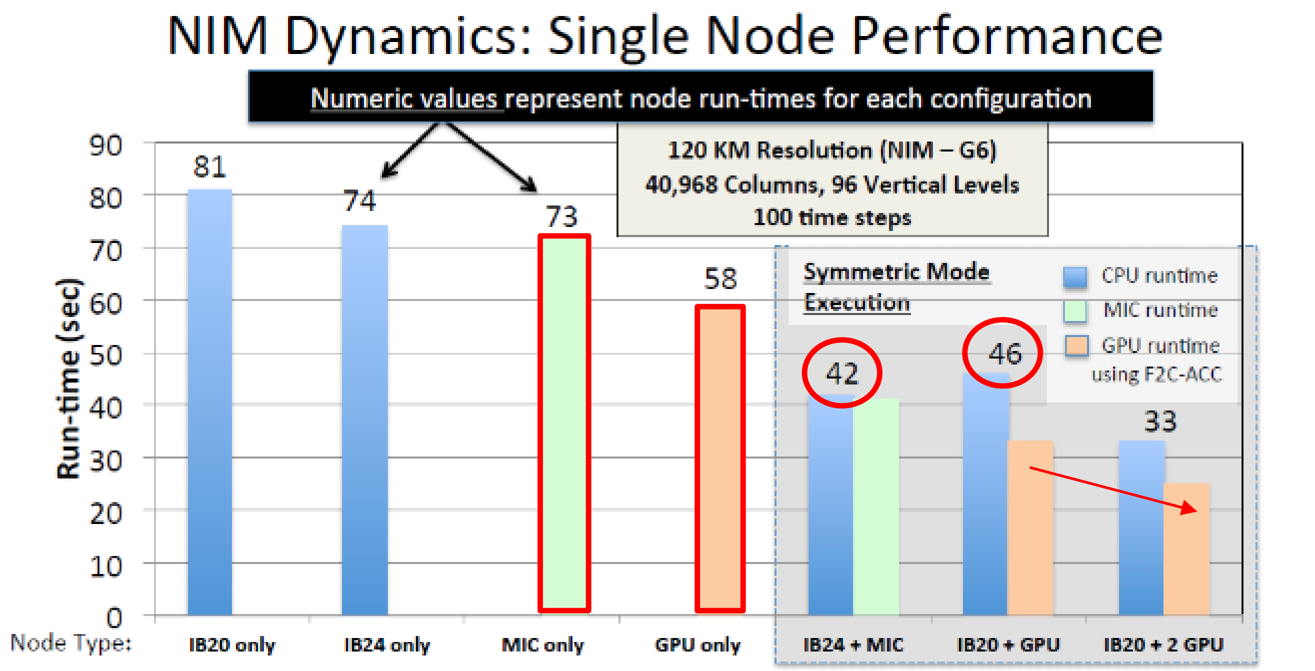

Parallelization and Performance

- Single source code (NIM rev 2724)
- Directive-based parallelization
  - OpenMP — CPU, MIC
  - F2C-ACC — GPU
  - SMS — MPI
  - OpenACC — GPU

System / Node configurations

- NVIDIA PSG Cluster
  - IB20: Intel IvyBridge, 20 cores, 3.0 GHz (Intel E5-2690 v2)
  - GPU: Kepler K40 2880 cores, 745 MHz, 12GB memory
- Intel Endeavor Cluster
  - IB24: Intel IvyBridge 24 cores, 2.70 GHz (Intel E5-2697 v2)
  - MIC: KNC 7120 61 cores, 1.238 GHz, 16 GB memory

http://www.esrl.noaa.gov/gsd/ab/ac/NIM-Performance.html

SUMMARY

K40 Speed-ups

2 x Ivy Bridge	1.3x
Xeon Phi	1.3x

NOTES:

- Results for standalone CPUs and accelerators
- K40 results from use of F2C-ACC

3. Tesla K40 Results: NOAA-NCEP and RRTMG

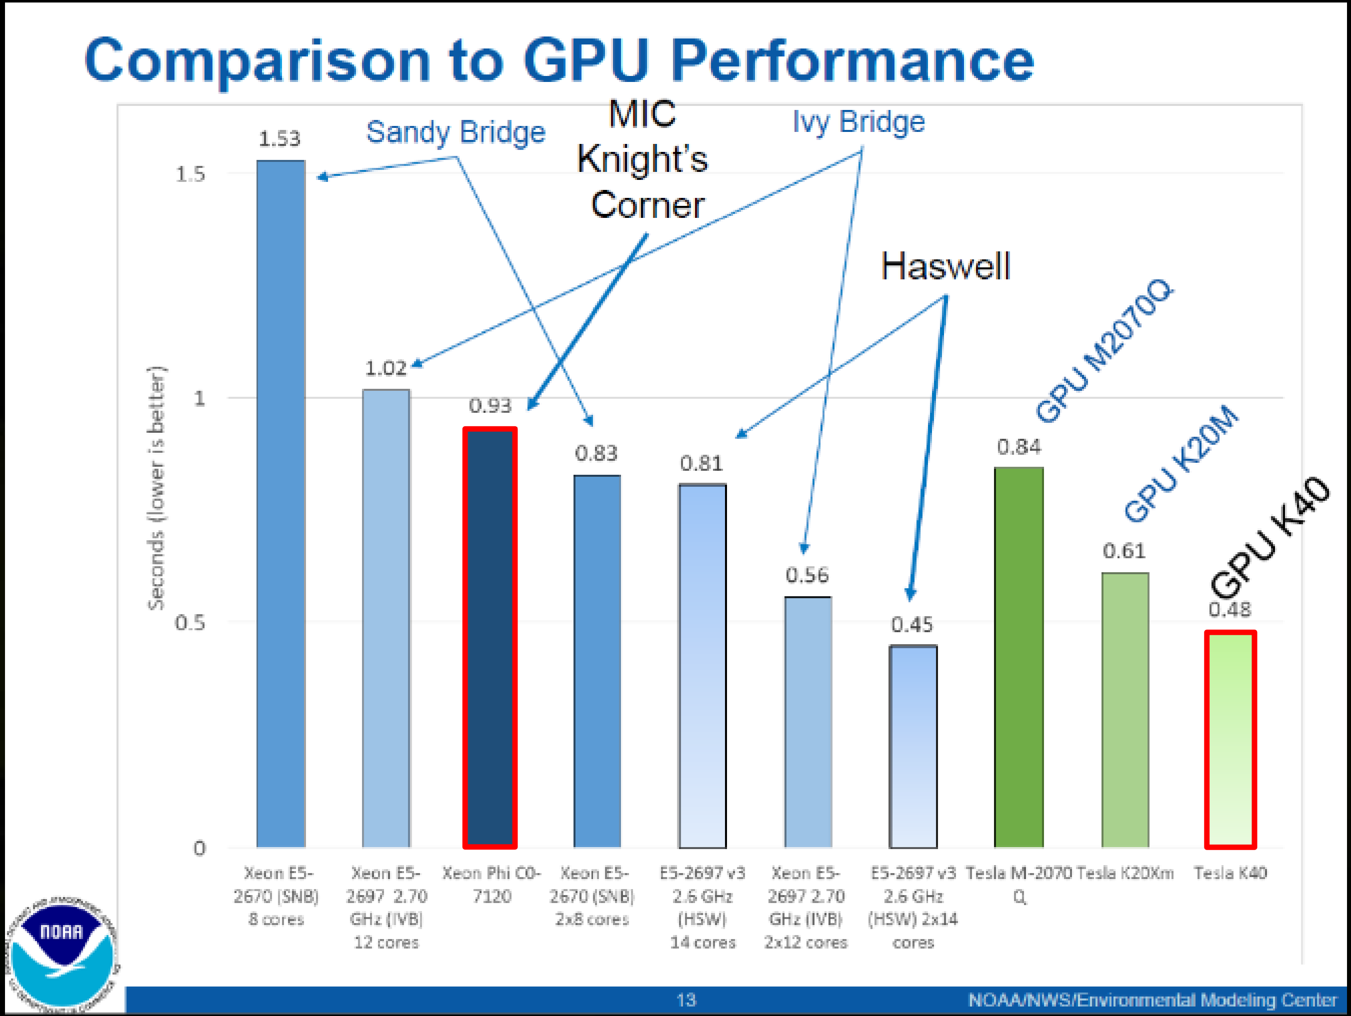

SUMMARY

K40 Speed-ups

2 x Sandy Bridge	1.7x
2 x Ivy Bridge	1.2x
2 x Haswell	0.9x
1 x Haswell	1.7x
Xeon Phi (Native)	1.9x

4. Tesla K40 Results: SSEC and WRF WSM5

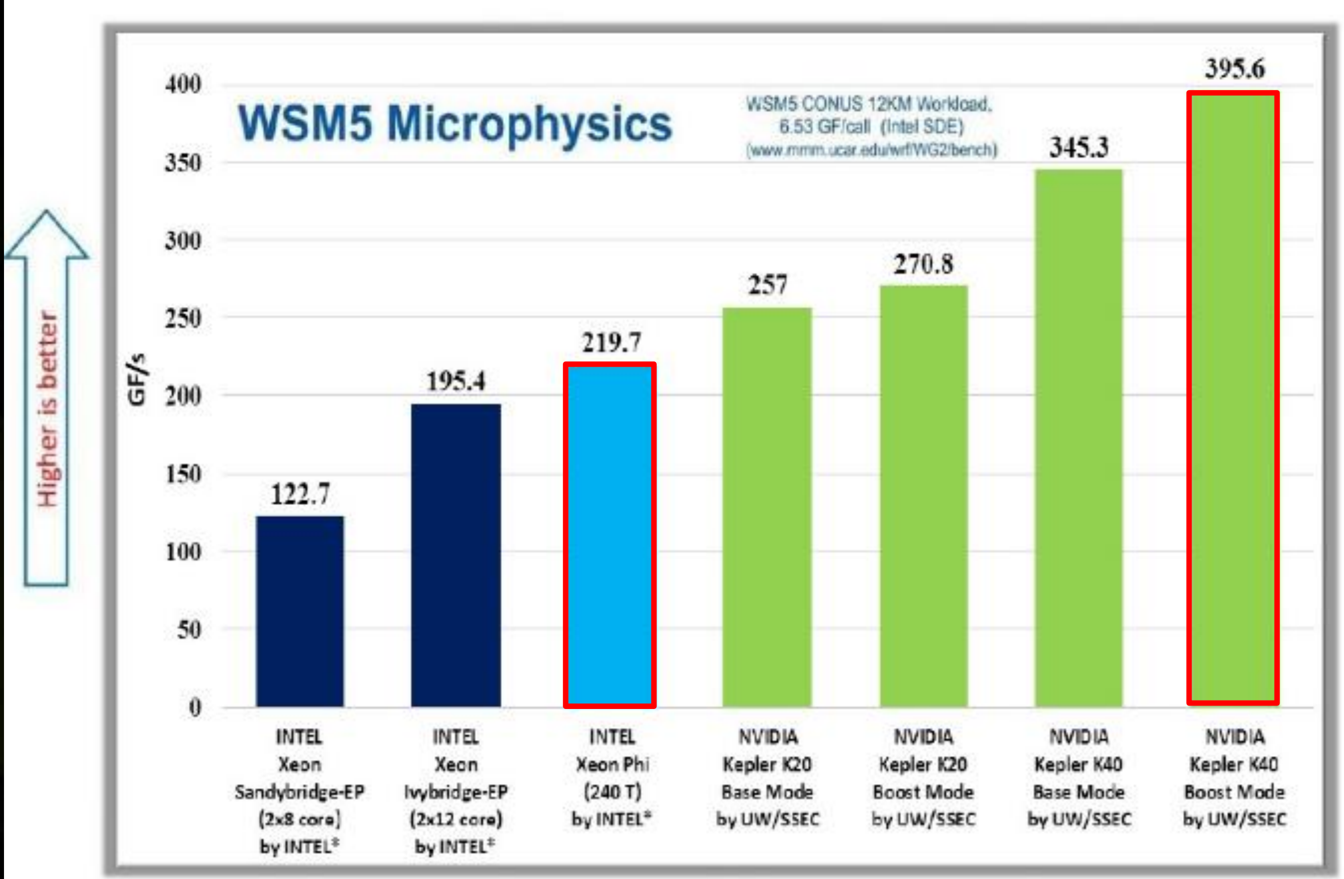

* *Code Restructuring to Improve Performance in WRF Model Physics on Intel Xeon Phi.* J. Michalakes. Workshop on Programming Weather, Climate and Earth System Models on Heterogeneous Multi-core Platforms, Boulder, Colorado, Sept. 19-20, 2013. (http://data1.gfdl.noaa.gov/multi-core/presentations/michalakes_5.pdf)

SUMMARY

K40 Speed-ups

2 x Sandy Bridge	3.2x
2 x Ivy Bridge	2.0x
Xeon Phi	1.8x

Tesla K40 Results from Multi-Core 4 Workshop

NVIDIA

NCAR UCAR CISL Computational & Information Systems Lab

2014 Heterogeneous Multi-Core 4 Workshop

September 17-18, 2014 at the National Center for Atmospheric Research in Boulder, Colorado

https://www2.cisl.ucar.edu/heterogeneous-multi-core-4-workshop/2014

Summary

- **Tesla K40 GPU the faster accelerator for each of the comparisons provided**

Workshop Study	NCAR CISL CESM DG Kernels	NOAA ESRL NIM	NOAA NCEP RRTMG	SSEC WRF WSM5
K40 / Phi Speed-up	1.6x	1.3x	1.9x	1.8x

- **IMPORTANT: all results are based on a CUDA programming environment**
  - **CUDA may not be suitable for production-level development/deployment**
  - **Most operational models require Fortran – therefore an OpenACC approach**

OpenACC Vital to ESM Production use of GPUs

NVIDIA Investments in OpenACC Standard

- NVIDIA acquisition of PGI, both OpenACC members; Ongoing Cray collaborations
- Support on important contributions from end-user members NOAA ESRL, CSCS/MCH, etc.
- Support/training to OpenACC members who have NWP/Climate GPU developments

OpenACC Workshops for the ESM Community

- Mar 2014: GTC OpenACC Roundtable for NWP and Climate Modeling (10 Models Reviewed)
- Sep 2014: NVIDIA Three Years HPC Roadmap Review for ESM (1/2 day OpenACC Focus)
- Oct 2014: DOE Oak Ridge Lab GPU and OpenACC HACK-A-THON (HYCOM, CICE, MPAS-O?)

  https://www.olcf.ornl.gov/training-event/hackathon-openacc2014/

Examples of OpenACC ESM Implementation

- Published OpenACC results on model success: COSMO, ICON, NICAM, NIM, NEMO

NVIDIA HPC Roadmap Review for ESM –Sep 2014

NVIDIA

19 Sep 14, Boulder, CO, USA – Following the Multi-Core 4 Workshop

Agenda:

Welcome and NVIDIA HPC Strategy Update for ESM Community

Roadmap of 3 Years Outlook for GPU Hardware and Software

- Dale Southard, Chief HPC Architect, NVIDIA Office of the CTO

Invited Talk on CUDA WRF Development Project at SSEC

- Dr. Jarno Mielikainen, Senior Staff Scientist, SSEC UW-Madison

OpenACC Features and Roadmap ½ Day Focus

OpenACC Features and Fortran Programming Considerations

- Jeff Larkin, Sr. HPC Engineer, NVIDIA Developer Technology Group

Cray Roadmap for OpenACC 2.0 and Beyond

- Eric Dolven, Cray

PGI Roadmap for OpenACC 2.0 and Beyond

- Dave Norton, Sr. HPC Applications Engineer, PGI

Delegates:

USA – NCAR (CISL, MMM); NOAA (ESRL, GFDL, NCEP); NASA GSFC; DOE (ORNL, ANL, PNNL, SNL)
INTL – Met Office (UK), STFC Daresbury (UK), KISTI (KR)

OpenACC Roundtable for NWP/Climate —Mar 2014

GPU TECHNOLOGY CONFERENCE

Model Contributions at GTC 2014 OpenACC Roundtable

- Motivation to identify critical and common OpenACC requests for international selection of 10 models

Model	Representatives
1. ASUCA	*Takashi Shimokawabe*, TiTech; *Michel Müller*, RIKEN
2. CAM-SE	*Jeff Larkin*, NVIDIA US; Matt Norman, ORNL
3. COSMO	*Peter Messmer*, NVIDIA CH; Claudio Gheller, Will Sawyer, CSCS
4. FIM/NIM	Mark Govett, NOAA
5. HARMONIE	*JC Desplat*, Enda O'Brien, ICHEC
6. ICON	*Peter Messmer*, NVIDIA CH; Claudio Gheller, Will Sawyer, CSCS
7. NEMO	Jeremy Appleyard, NVIDIA UK
8. NICAM	*Akira Naruse*, NVIDIA JP; Hisashi Yashiro, RIKEN
9. WRF	Carl Ponder, NVIDIA US
10. COAMPS	*Dave Norton*, PGI; Gopal Patnaik, US NRL

Results of GTC OpenACC Roundtable (I)

OpenACC Feature Request	PGI Response
Support for derived types with member arrays in data clauses	Initial support in 14.4 released Apr 2014
Deep copy support: • Arrays of derived type with member arrays of derived type with member arrays, etc.; CAM-SE, COSMO, ICON, HARMONIE • Deep and shallow copy for nested derived types, selective deep copy; ICON, HARMONIE	Not supported, plans not yet announced
Minimize code restructuring such as inner-outer loop reordering; CAM-SE, others	In development, initial support in 14.7 (Jul)
Want consistent experiences across vendor implementations: • Perhaps more strict specification so different implementations have less freedom to interpret; NIM • User desires to write program once that will port to different systems/compilers	No response, but close collaboration with OpenACC member alliance on specifications

Results of GTC OpenACC Roundtable (II)

GPU TECHNOLOGY CONFERENCE

NVIDIA

OpenACC Feature Request	PGI Response
Allow allocation of noncontiguous array sections such as matrix interior; HARMONIE	Not supported, plans not yet announced
Allow unallocated arrays in data clauses; HARMONIE	In development, initial support in 14.7 (Jul)
Bitwise consistency between CPU and GPU; ICON	Investigating for support, no release plan
Expose pinned memory; NIM, NICAM	No plans announced, but under consideration
Expose shared memory; WRF	No plans announced, but under consideration
Reorder array dimensions for CPU vs. GPU; ASUKA, others	Not supported, plans not yet announced
Full procedure support with no restrictions, and no specification required on subroutine definition; COAMPS	Not supported, plans not yet announced
Expose managed memory; NEMO	Investigating with CUDA 6, support in 14.7
CUDA Aware MPI or OpenACC aware MPI	No plans announced, but under consideration

COSMO

Towards GPU-accelerated Operational Weather Forecasting

- Oliver Fuhrer (MeteoSwiss), NVIDIA GTC 2013, Mar 2013

Source: http://on-demand.gputechconf.com/gtc/2013/presentations/S3417-GPU-Accelerated-Operational-Weather-Forecasting.pdf

Implementation of COSMO on Accelerators

- Oliver Fuhrer (MeteoSwiss), ECMWF Scalability Workshop, Apr 2014

Source: http://old.ecmwf.int/newsevents/meetings/workshops/2014/Scalability/

Piz Daint System http://www.cscs.ch/piz_daint/

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

- Piz Daint consists of 5272 compute nodes with 169 TB of total system memory
- Each node contains 1 x Intel Xeon E5-2670 CPU and 1 x NVIDIA K20X GPU
- COSMO operational tests up to 1000's of nodes on Piz Daint Cray XC30

COSMO Results from ECMWF Scalability Workshop

Implementation of COSMO on Accelerators

-by Dr. Oliver Fuhrer, MeteoSwiss; ECMWF Scalability Workshop, Apr 2014

CUDA-based DSL "STELLA" for dynamics + OpenACC for physics and data assimilation

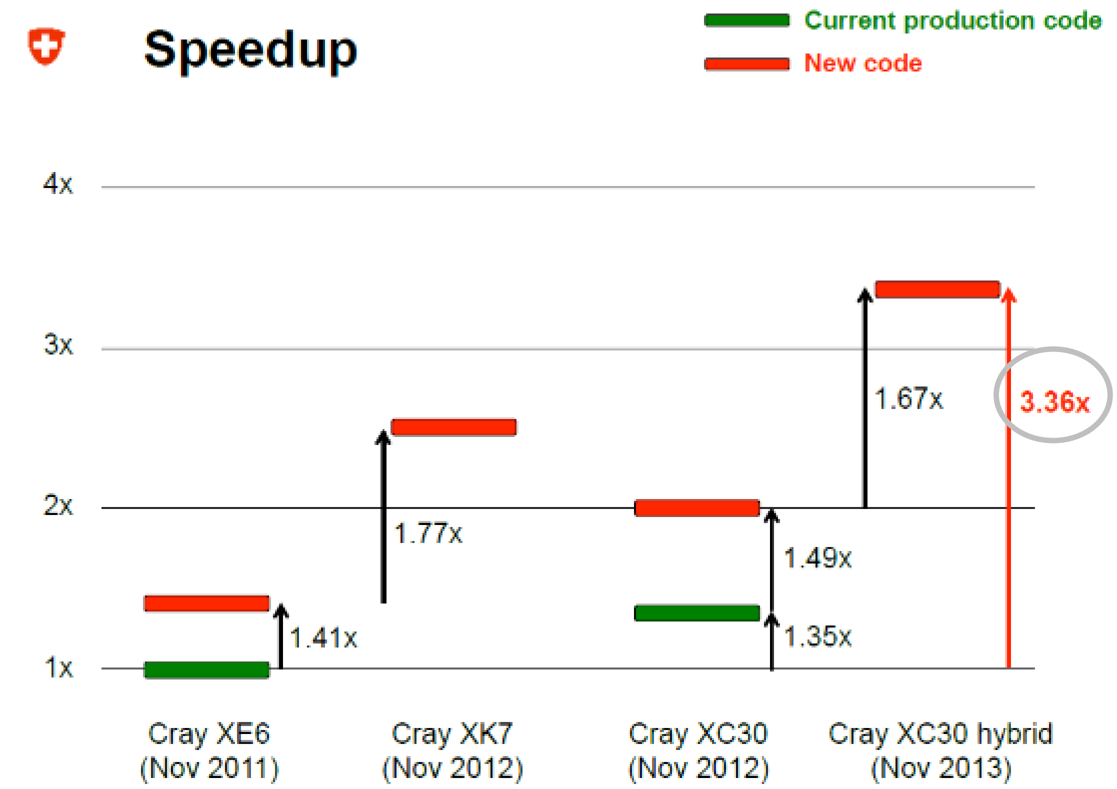

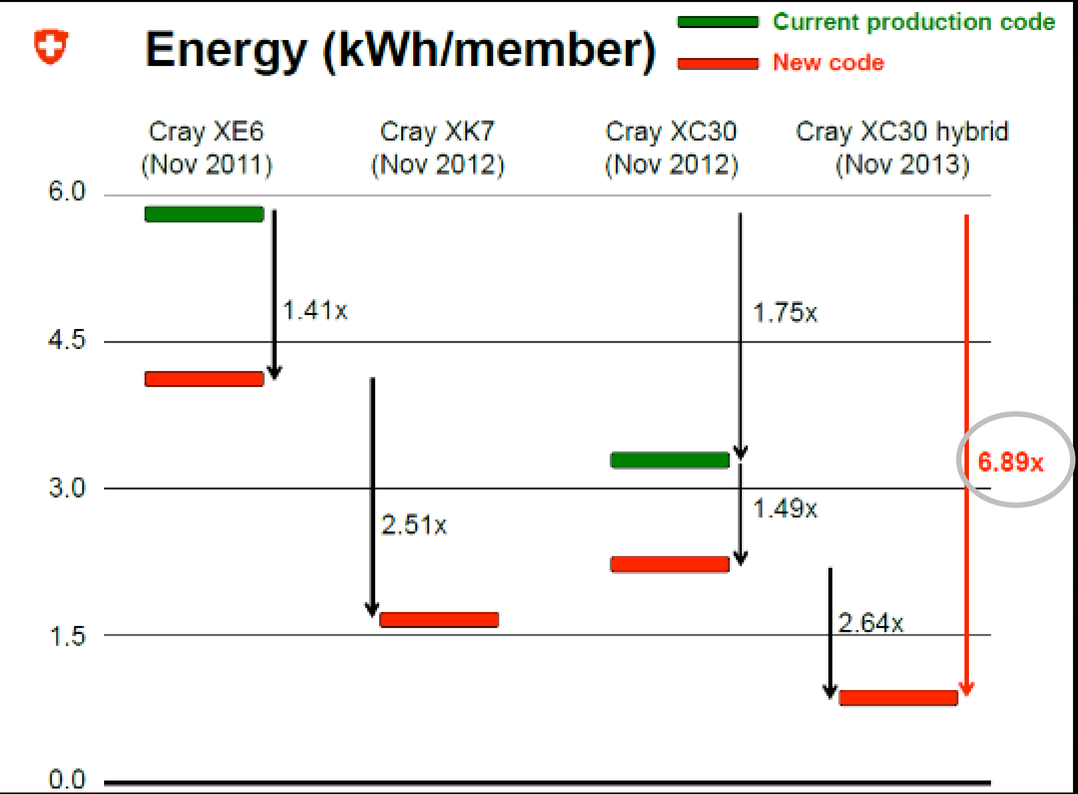

Source: http://old.ecmwf.int/newsevents/meetings/workshops/2014/Scalability/Presentations/pdfs/Fuhrer.pdf

NEMO

Accelerating NEMO with OpenACC

- Maxim Milakov (NVIDIA), NVIDIA GTC 2013, Mar 2013

Source: http://on-demand.gputechconf.com/gtc/2013/presentations/S3209-Accelerating-NEMO-with-OpenACC.pdf

NEMO on GPU-based Heterogeneous Architectures: a Case Study Using OpenACC

- Jeremy Appleyard (NVIDIA), NEMO UGM, Jul 2014

NEMO Performance with OpenACC and GPUs

NEMO Model

http://www.nemo-ocean.eu/

- Nucleus for European Modelling of the Ocean global and regional OGCM
- Primary developers CNRS, Mercato-Ocean, UKMO, NERC, CMCC, INGV
- OCN component for 5 of 7 Earth system models in the ENES http://enes.org
- European consortium of 40 projects, 400 users, and ~50 publications/year

Configurations

- GYRE50: Idealized double gyres, 1/4° horizontal resolution, 31 vertical layers
- ORCA025: Global high resolution, 1/4° horizontal resolution, 75 vertical layers

NVIDIA "PSG" Cluster

http://psgcluster.nvidia.com/trac

- PSG consists of 30 compute nodes of mixed type, each 128 GB of system memory
- This study: Each node 2 x Intel Xeon Ivy Bridge CPUs and 6 x NVIDIA K40 GPUs
- NEMO tests on 8 nodes using 20 of 20 cores per node, and 2 of 6 GPUs per node

NVIDIA PSG Cluster

NEMO Coupling to European Climate Models

- NEMO critical for European climate models: ocean component for 5 of 7 modeling groups*

ENES
European Network for Earth System modelling
http://enes.org

Country	name of model (CMIP5)	Atmosphere	Ocean	Sea Ice	Coupler	Land Surface *Vegetation	Atmospheric Chemistry	Ocean Bio-geochemistry
Consortium	EC-EARTH	IFS	NEMO	LIM	OASIS	HTESSEL	TM5	
France	IPSLCM5	LMDz	NEMO	LIM	OASIS	ORCHIDEE	INCA	PISCES
France	CNRM-Cerfacs	ARPEGE	NEMO	GELATO	OASIS	SURFEX		
Germany	MPI-ESM	ECHAM5	MPIOM	MPIOM	OASIS	JSBACH*	HAM	HAMOCC
Italy	C-ESM	ECHAM5	NEMO	LIM	OASIS	SILVA		PELAGOS
UK	HadGEM2	UM	UM	CICE	OASIS	TRIFFID*	UKCA	diat-HADOCC
Norway	NorESM	NCAR	MICOM	CICE	CPL7	CLM	Chemistry	HAMOCC
EC-Earth Con	Netherlands, Sweden, Ireland, Denmark, Spain, Portugal, Italy, Belgium							

* UKMO has announced NEMO as the ocean component model for HadGEM3

NEMO Performance with OpenACC and GPUs

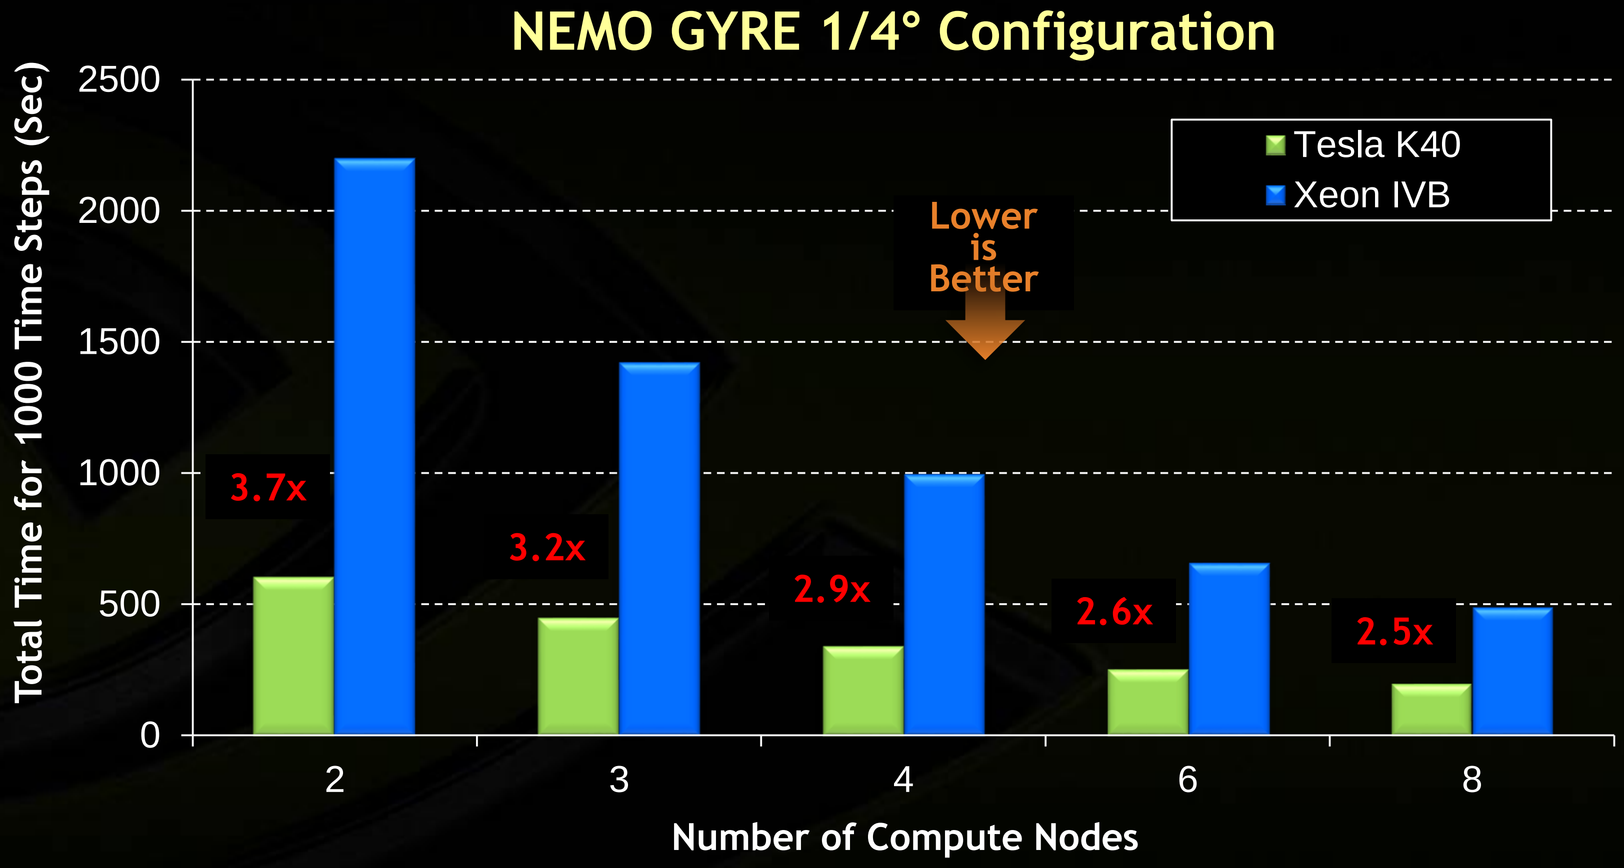

PSG node utilization:
2 x IVB + 2 x K40

Without using GPUs

Use of GPUs

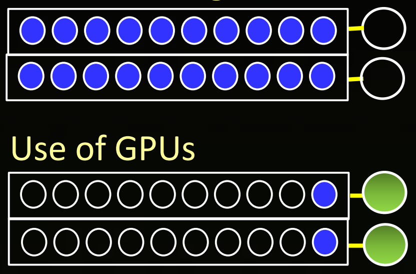

GYRE settings:

- **Output every 5 days**
- **Time steps = 1000**
- **NEMO release 3.5**

NEMO and GPU Performance with OpenACC

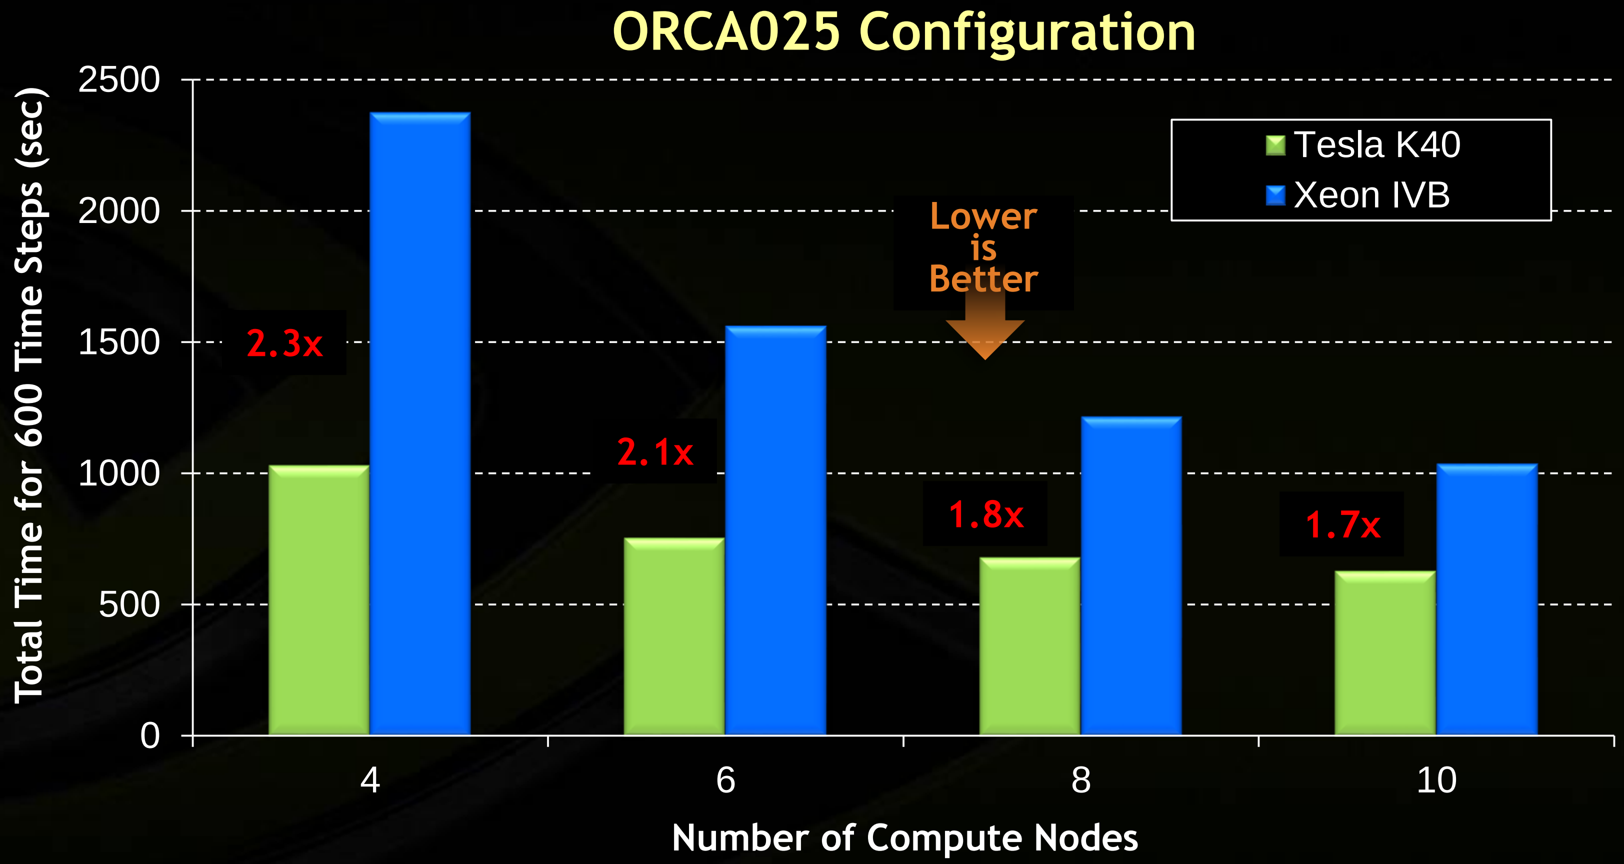

Node utilization:
2 x IVB + 2 x K40

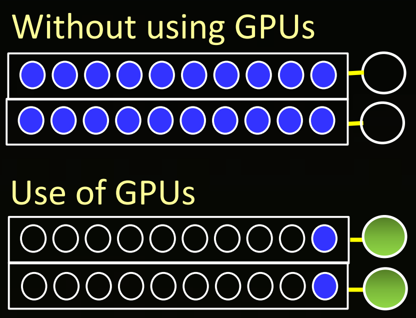

ORCA025 settings:

- Output every 5 days
- Total run: 10 days
- Time steps: 600
- NEMO 3.5

NEMO and GPU Performance with OpenACC

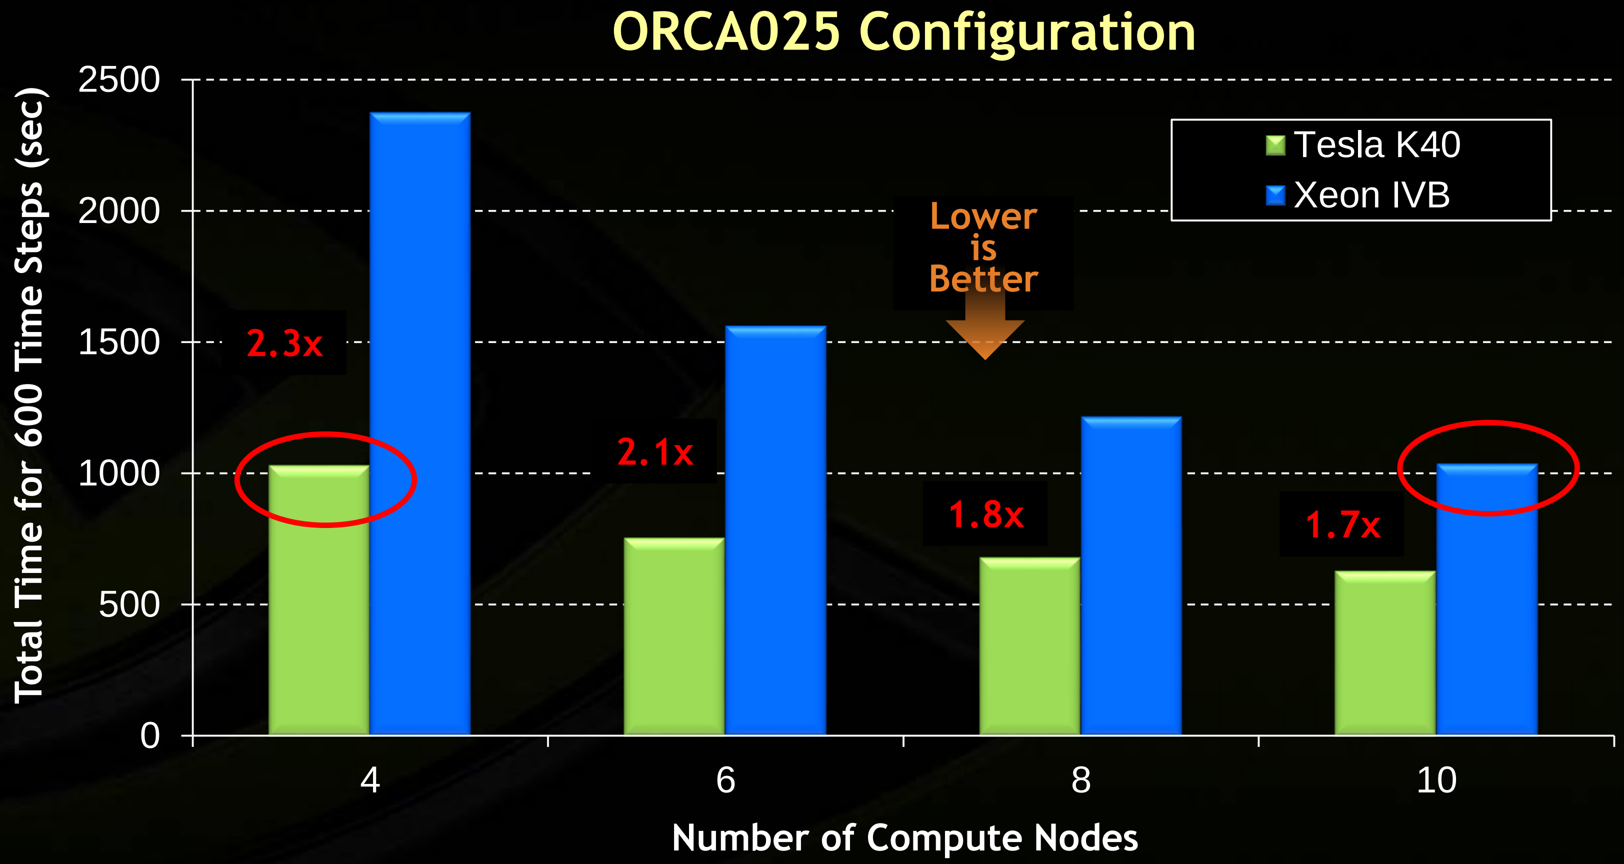

Node utilization:
2 x IVB + 2 x K40

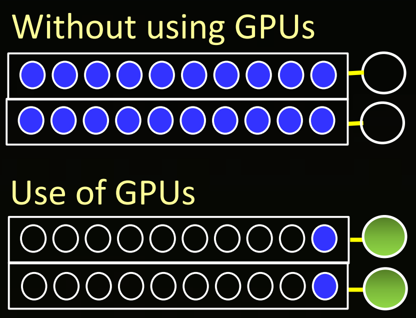

ORCA025 settings:

- Output every 5 days
- Total run: 10 days
- Time steps: 600
- NEMO 3.5

NEMO HPC Configurations at Equal Performance

2 nodes + 8 GPUs = 4 nodes + 8 GPUs = 10 nodes

- Flexibility: GPUs free-up existing HPC nodes/cores for other applications
- Efficiency: GPU-based nodes more cost effective for new HPC purchase

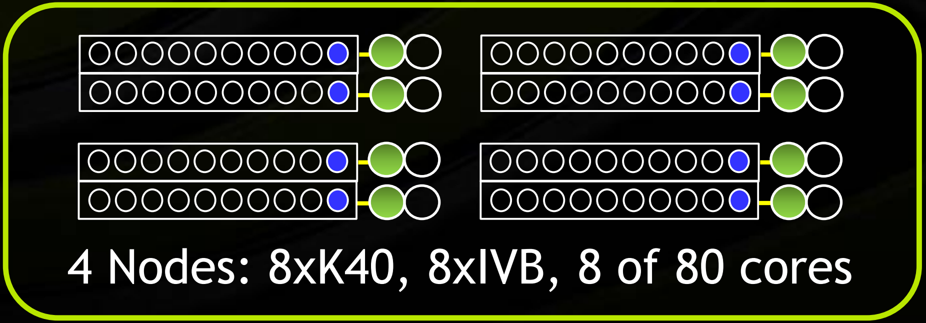

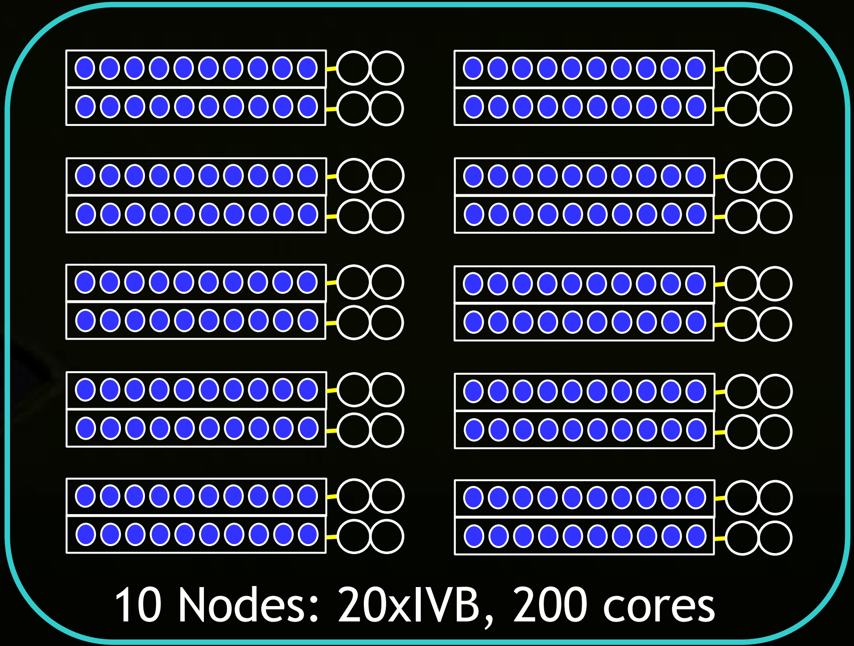

NICAM

Recent Performance of NICAM on the K-Computer and Activities Towards Post-Petascale Computing

- Hisashi Yashiro(RIKEN), ECMWF Scalability Workshop, Apr 2014

Source: http://www.ecmwf.int/newsevents/meetings/workshops/2014/Scalability/Presentations/pdfs/Yashiro.pdf

GPU optimization of global cloud resolving model NICAM core dynamics using OpenACC

- Hisashi Yashiro(RIKEN), GTC Japan, Jul 2014

Source: http://www.gputechconf.jp/page/sessions.html

NICAM Performance with OpenACC and GPUs

NICAM Model http://nicam.jp

- Nonhydrostatic ICosahedral Atmospheric Model (2000 - Tomita and Satoh)
- Primary developers are JAMSTEC, University of Tokyo, and RIKEN AICS
- Collaborations: Athena project, COLA, ECMWF, NICS/UTK, ICOMEX
- Global resolution has high as 1.75 km achieved on K computer in Japan

NICAM-DC Project http://scale.aics.riken.jp/nicamdc/

- Dynamical core package of full NICAM full model - available as open source
- High order (3rd) advection scheme based on conservative semi-Lagrangian

TSUBAME 2.5 System http://www.gsic.titech.ac.jp/en/tsubame

- TSUBAME 2.5 consists of 1408 compute nodes with 52 GB of system memory
- Each node 2 x Intel Xeon Westmere-EP 2.9 GHz CPUs and 3 x NVIDIA K20X GPUs
- NICAM tests on 5 nodes using 8 of 12 cores per node, and 2 of 3 GPUs per node

東京工業大学 Tokyo Institute of Technology

NICAM Performance with OpenACC and GPUs

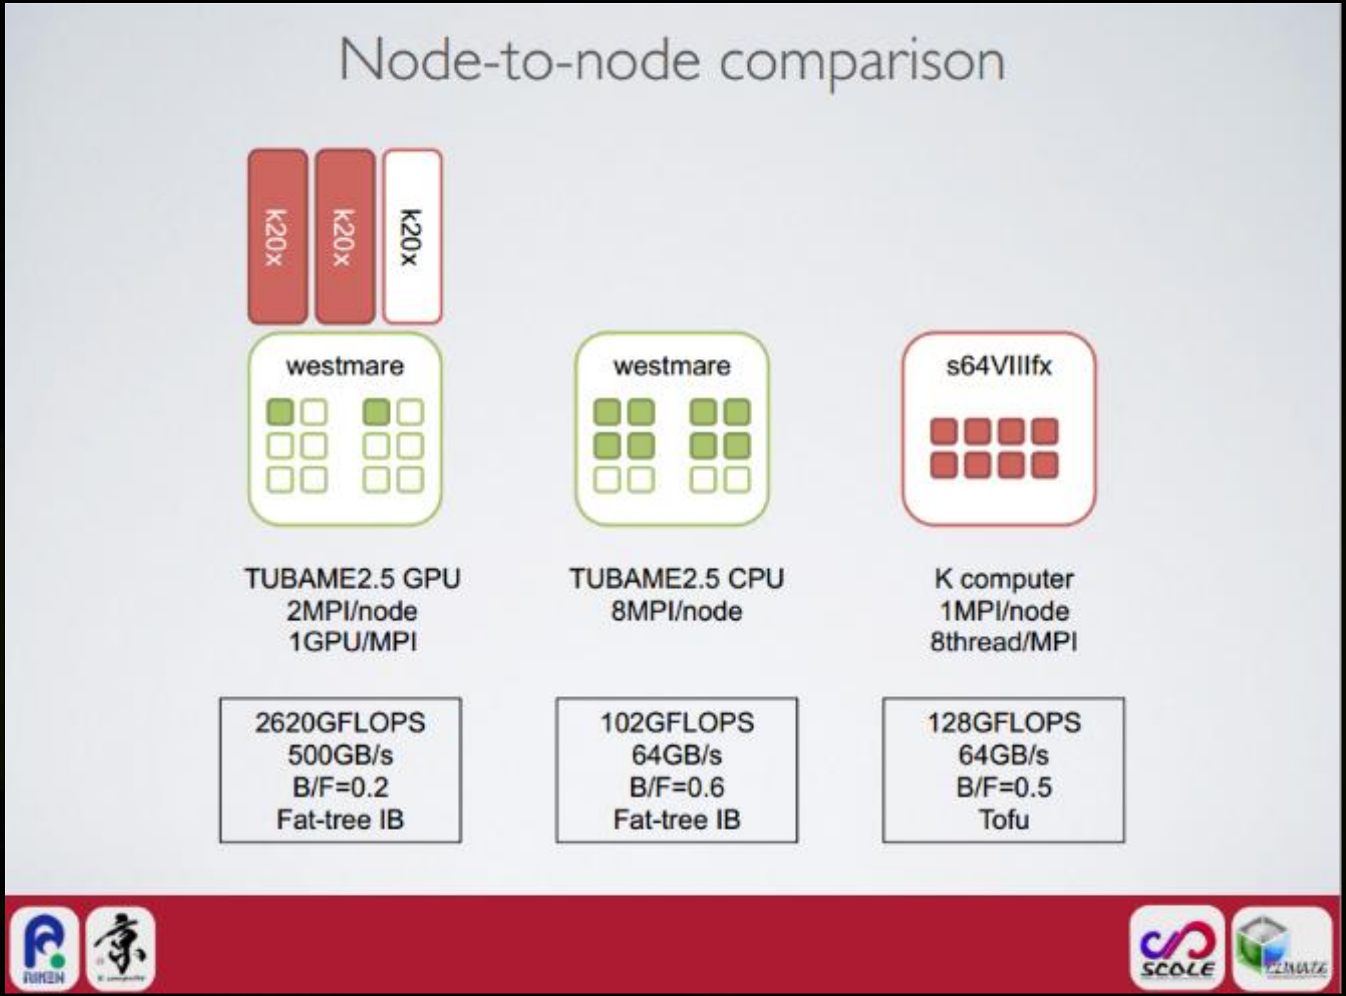

From the Workshop:
ECMWF Scalability Project
Apr 2014, Reading, UK

Recent Performance of NICAM on the K-Computer and Activities Towards Post-Petascale Computing
-by Dr. Hisashi Yashiro, RIKEN AICS

TSUBAME 2.5
Use of 2 cores + 2 GPUs
vs. 8 of 12 cores per node

Source: http://www.ecmwf.int/newsevents/meetings/workshops/2014/Scalability/Presentations/pdfs/Yashiro.pdf

NICAM Performance with OpenACC and GPUs

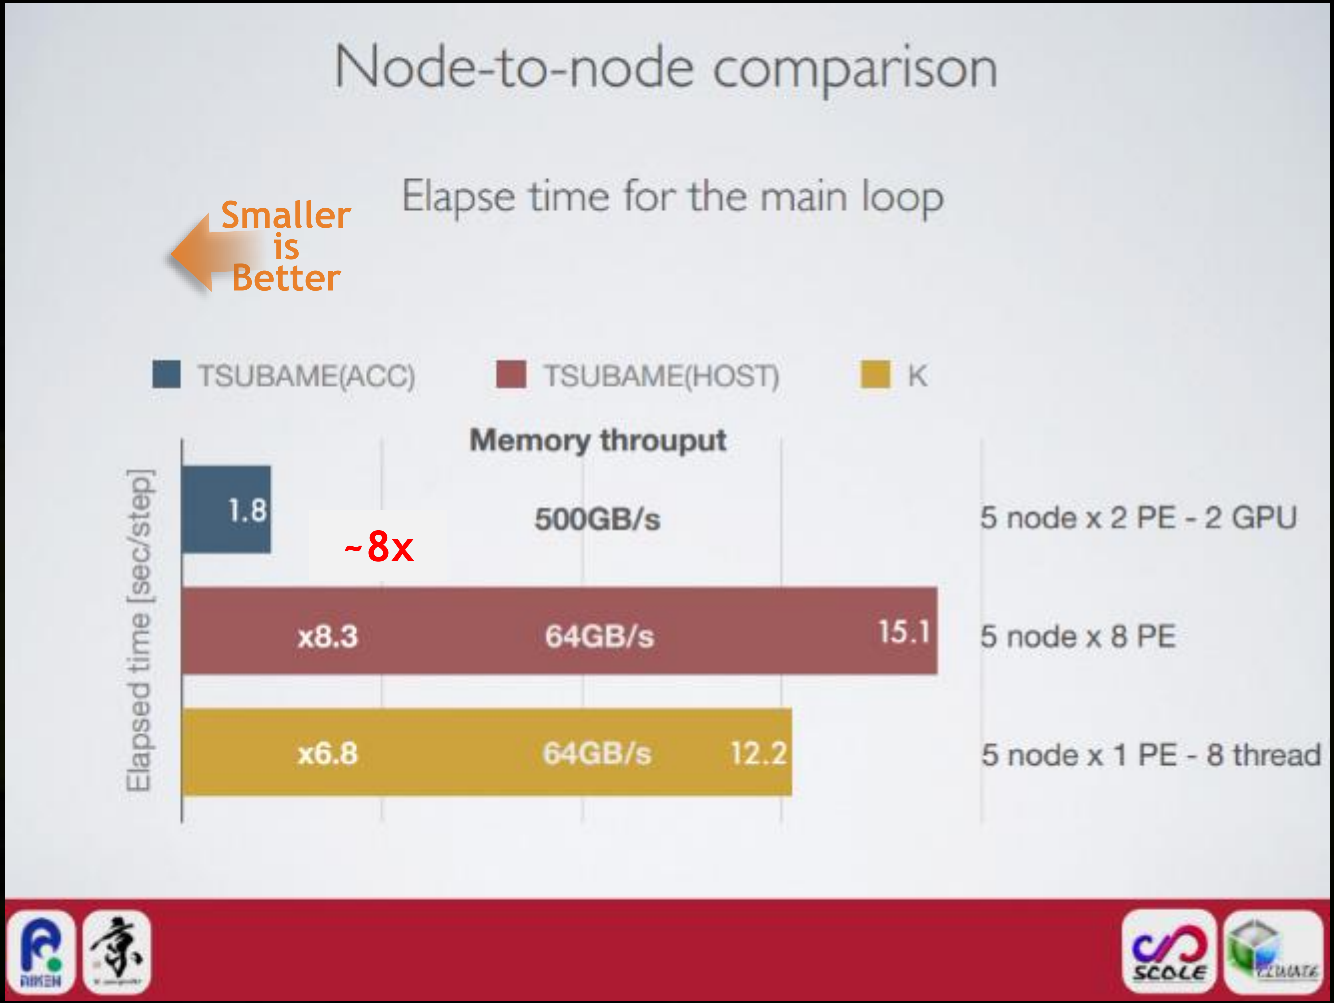

From the Workshop:
ECMWF Scalability Project
Apr 2014, Reading, UK

Recent Performance of NICAM on the K-Computer and Activities Towards Post-Petascale Computing
-by Dr. Hisashi Yashiro, RIKEN AICS

Observed >8x speedup for dynamics using OpenACC on 5 nodes with 10 x K20X vs. 10 x Intel CPU (40 cores)

Source: http://www.ecmwf.int/newsevents/meetings/workshops/2014/Scalability/Presentations/pdfs/Yashiro.pdf

NICAM Performance with OpenACC and GPUs

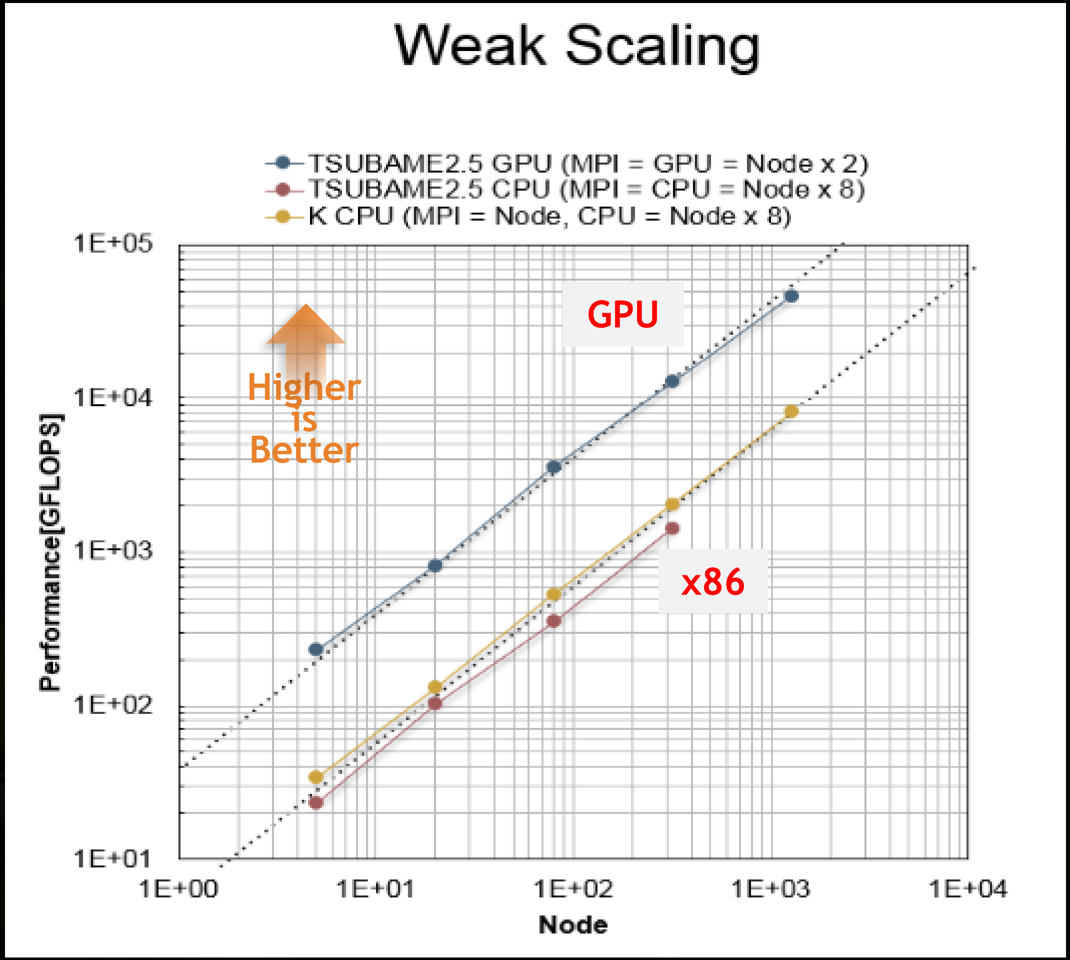

From the Workshop:
ECMWF Scalability Project
Apr 2014, Reading, UK

Recent Performance of NICAM on the K-Computer and Activities Towards Post-Petascale Computing
-by Dr. Hisashi Yashiro, RIKEN AICS

Observed GPU scaling to >1000 nodes (>2000 GPUs)

Source: http://www.ecmwf.int/newsevents/meetings/workshops/2014/Scalability/Presentations/pdfs/Yashiro.pdf

Agenda: NVIDIA HPC Directions for ESM

- NVIDIA HPC and GPU Update
- ES Model Progress on GPUs
- **GPU Technology Roadmap**

New CPU Platforms Available During 2014

Tesla K40 (Kepler)

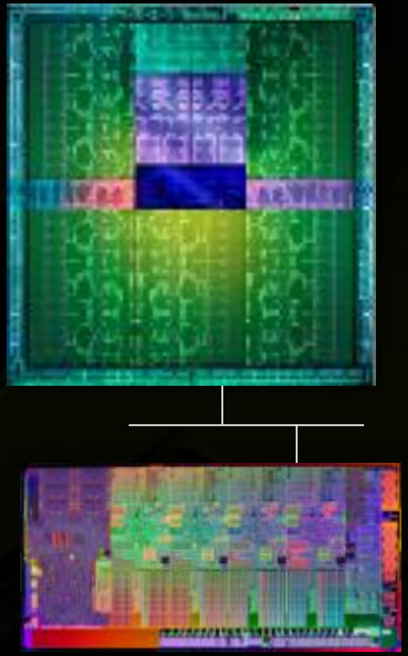

PCIe - 16 GB/s

X86 | ARM64 | Power8

Choice of CPU Platforms Available Starting 2016

Pascal

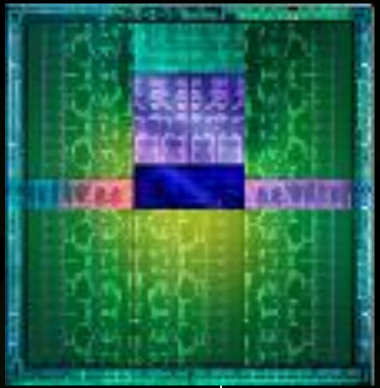

PCIe - 16 GB/s

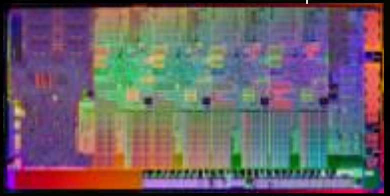

X86 | ARM64 | Power

Pascal

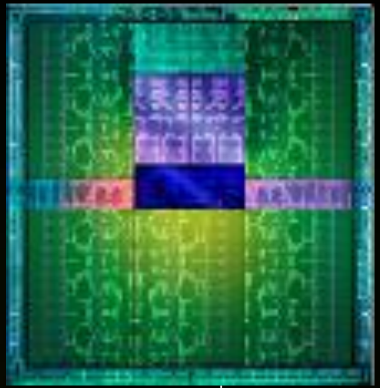

NVLink - 80 GB/s

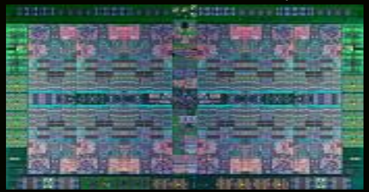

ARM64 | Power

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)

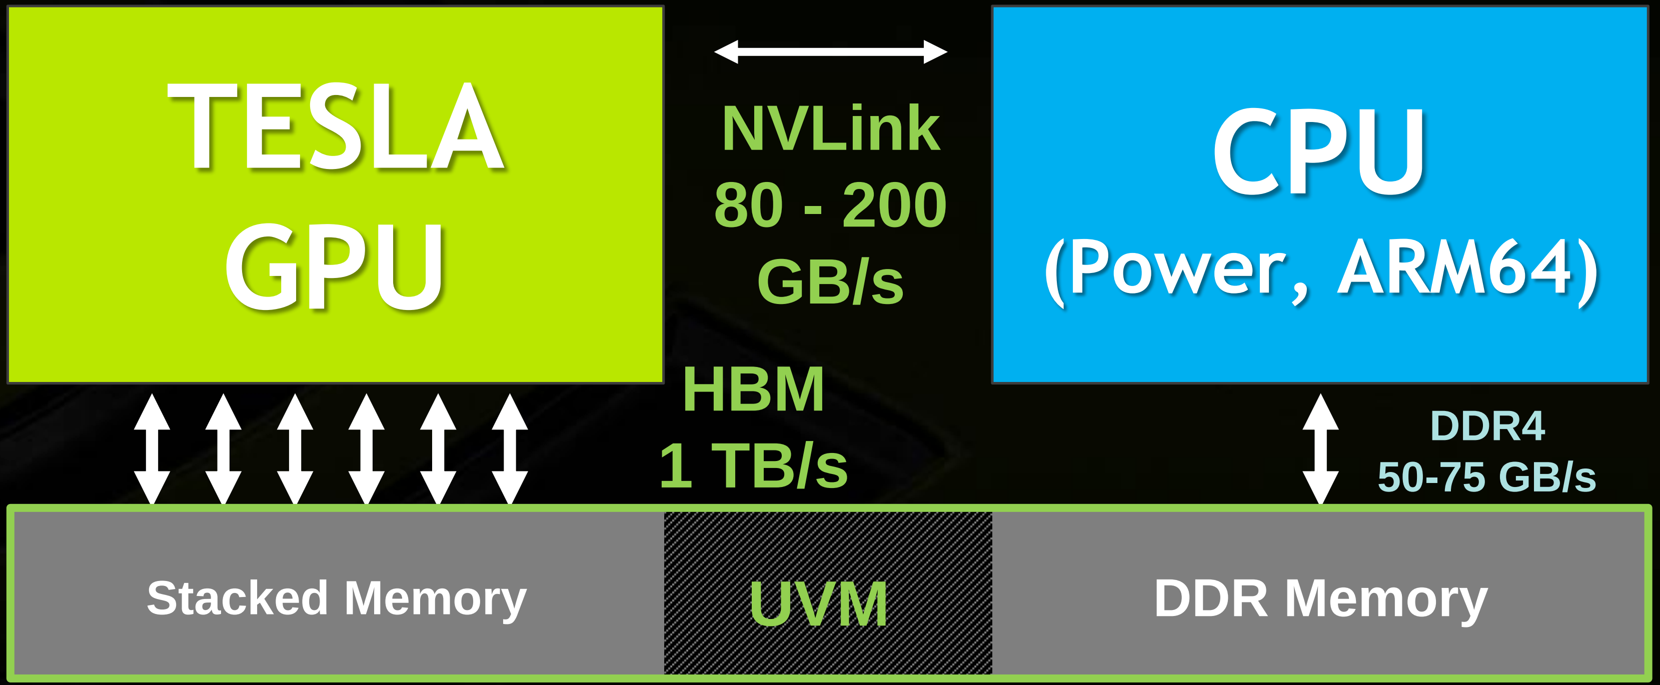

Programming Strategies for GPU Acceleration

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

Increasing Development Effort →

NVIDIA AmgX for Iterative Implicit Methods

- Scalable linear solver library for Ax = b iterative methods
- No CUDA experience required, C API: links with Fortran, C, C++
- Reads common matrix formats (CSR, COO, MM)
- Interoperates easily with MPI, OpenMP, and hybrid parallel
- Single and double precision; Supported on Linux, Win64
- Multigrid; Krylov: GMRES, PCG, BiCGStab; Preconditioned variants
- Classic Iterative: Block-Jacobi, Gauss-Seidel, ILU's; Multi-coloring
- Flexibility: All methods as solvers, preconditioners, or smoothers
- Download AmgX library: http://developer.nvidia.com/amgx

NVIDIA AmgX Weak Scaling on Titan 512 GPUs

Use of 512 nodes on ORNL TITAN System

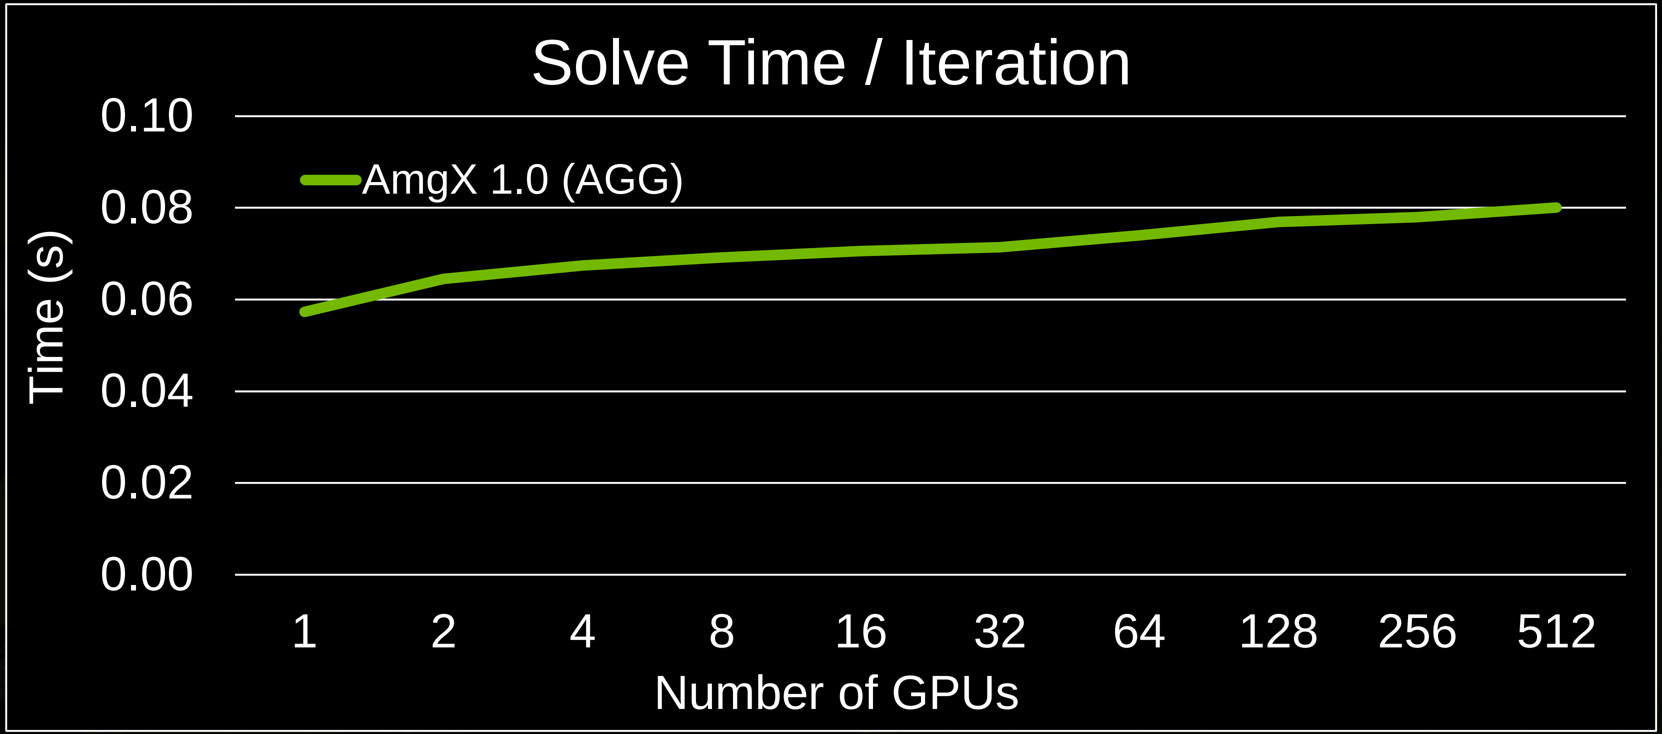

- **Poisson matrix with ~8.2B rows solved in under 13 sec (200e3 Poisson matrix per GPU)**
- **ORNL TITAN: NVIDIA K20X one per node; CPU 16 core AMD Opteron 6274 @2.2GHz**

ACME: US DOE Accelerator-Based Climate Model

- **ACME: Accelerated Climate Model for Energy**
  - 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)
  - Towards NH global atm 12 km, ocn 15 km, 80 years

- **Co-design project with US DOE LCF systems**

First Report – 11 Jul 2014

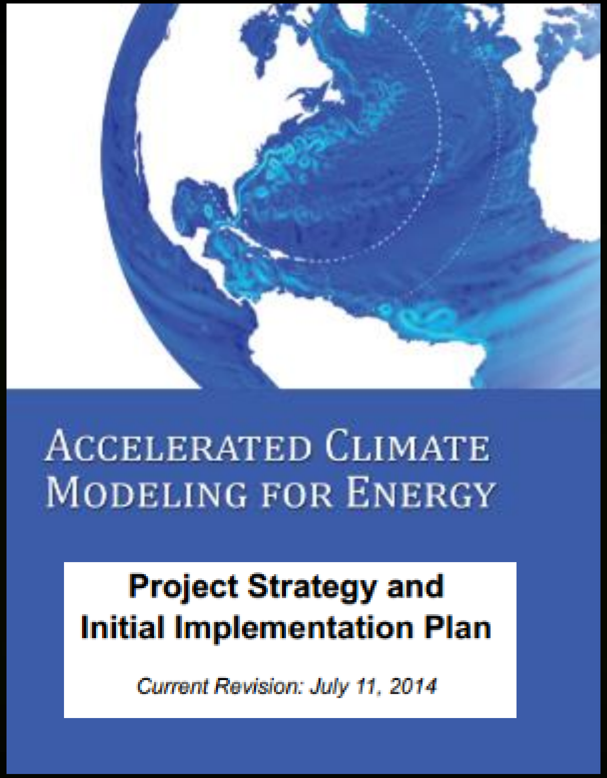

Source: http://climatemodeling.science.energy.gov/sites/default/files/publications/acme-project-strategy-plan_0.pdf

ACME: US DOE Accelerator-Based Climate Model

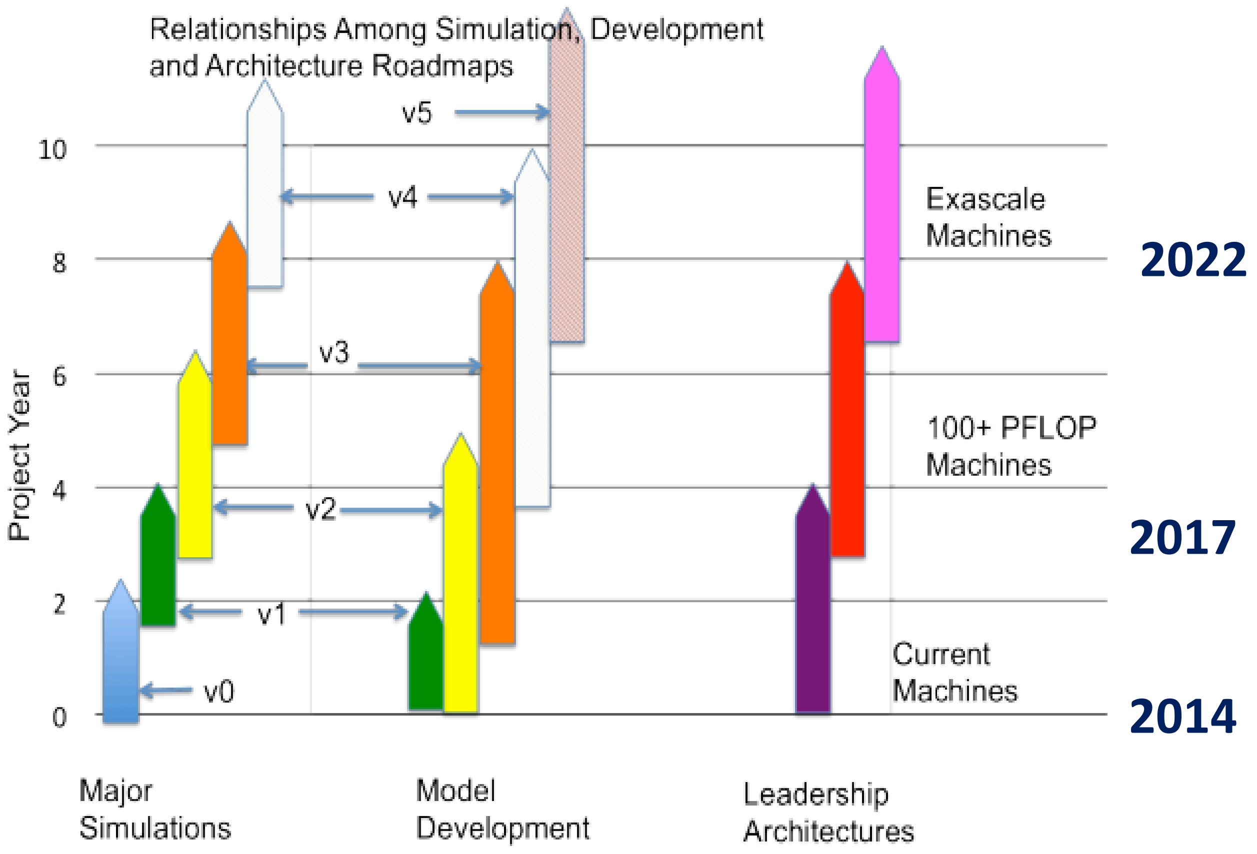

ACME Project Roadmap
(page 2 in report)

- **TBD**

- **CORAL: ORNL – ANL – LLNL**
  [Systems to be announced SC14]
- **Trinity, NERSC8 (Cori)**

- **TITAN - ORNL** [AMD + GPU]
- **Mira – ANL** [Blue GeneQ]

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

Summary Highlights: NVIDIA HPC Directions for ESM

- NVIDIA observes strong ESM community interest in GPU acceleration
  - New technologies: Pascal, NVLink, more CPU platform choices
  - NVIDIA business and engineering collaborations in 16 model projects
  - Investments in OpenACC: PGI release of 14.9; Continued Cray collaborations

- GPU progress for several models – we examined a few of these
  - MeteoSwiss developments towards COSMO operational NWP on GPUs
  - OpenACC for NICAM: 1 of 6 applications driving exascale plans at RIKEN, JP
  - OpenACC for NEMO: Leading ocean component model in ENES models
  - Ongoing GPU developments for 24 of 30 ES models where NVIDIA focus

- Watch for announcements coming at Supercompting'14, 17 Nov 2014
  - New GPU products and technologies with further details on roadmap

Thank you and Questions?

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