GPU Benefits for Earth System Science

Stan Posey, Program Manager, ESS Domain NVIDIA (HQ), Santa Clara, CA
Raghu Kumar, PhD, Software Engineer, Developer Technology NVIDIA, Boulder, CO

TOPICS

- ACHIEVEMENTS IN HPC AND AI
- NUMERICAL MODEL DEVELOPMENTS
- GPU UPDATE ON MPAS-A

World-Leading HPC Systems Deploy NVIDIA GPUs

ORNL Summit
#1 Top 500
27,648 GPUs| 144 PF

LLNL Sierra
#2 Top 500
17,280 GPUs| 95 PF

Piz Daint
Europe's Fastest
5,704 GPUs| 21 PF

ABCI
Japan's Fastest
4,352 GPUs| 20 PF

ENI HPC4
Fastest Industrial
3,200 GPUs| 12 PF

NERSC-9 HPC System Based on "Volta-Next" GPU During 2020:

Perlmutter: A System Optimized for Science

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

SC18 Gordon Bell Award: NERSC and NVIDIA Team

Exascale Deep Learning for Climate Analytics

Thorsten Kurth* tkurth@lbl.gov
Sean Treichler† sean@nvidia.com
Joshua Romero† joshr@nvidia.com
Mayur Mudigonda* mudigonda@berkeley.edu
Nathan Luehr† nluehr@nvidia.com
Everett Phillips† ephillips@nvidia.com
Ankur Mahesh* amahesh@lbl.gov
Michael Matheson‡ mathesonma@ornl.gov
Jack Deslippe* jrdeslippe@lbl.gov
Massimiliano Fatica† mfatica@nvidia.com
Prabhat* prabhat@lbl.gov
Michael Houston† mhouston@nvidia.com

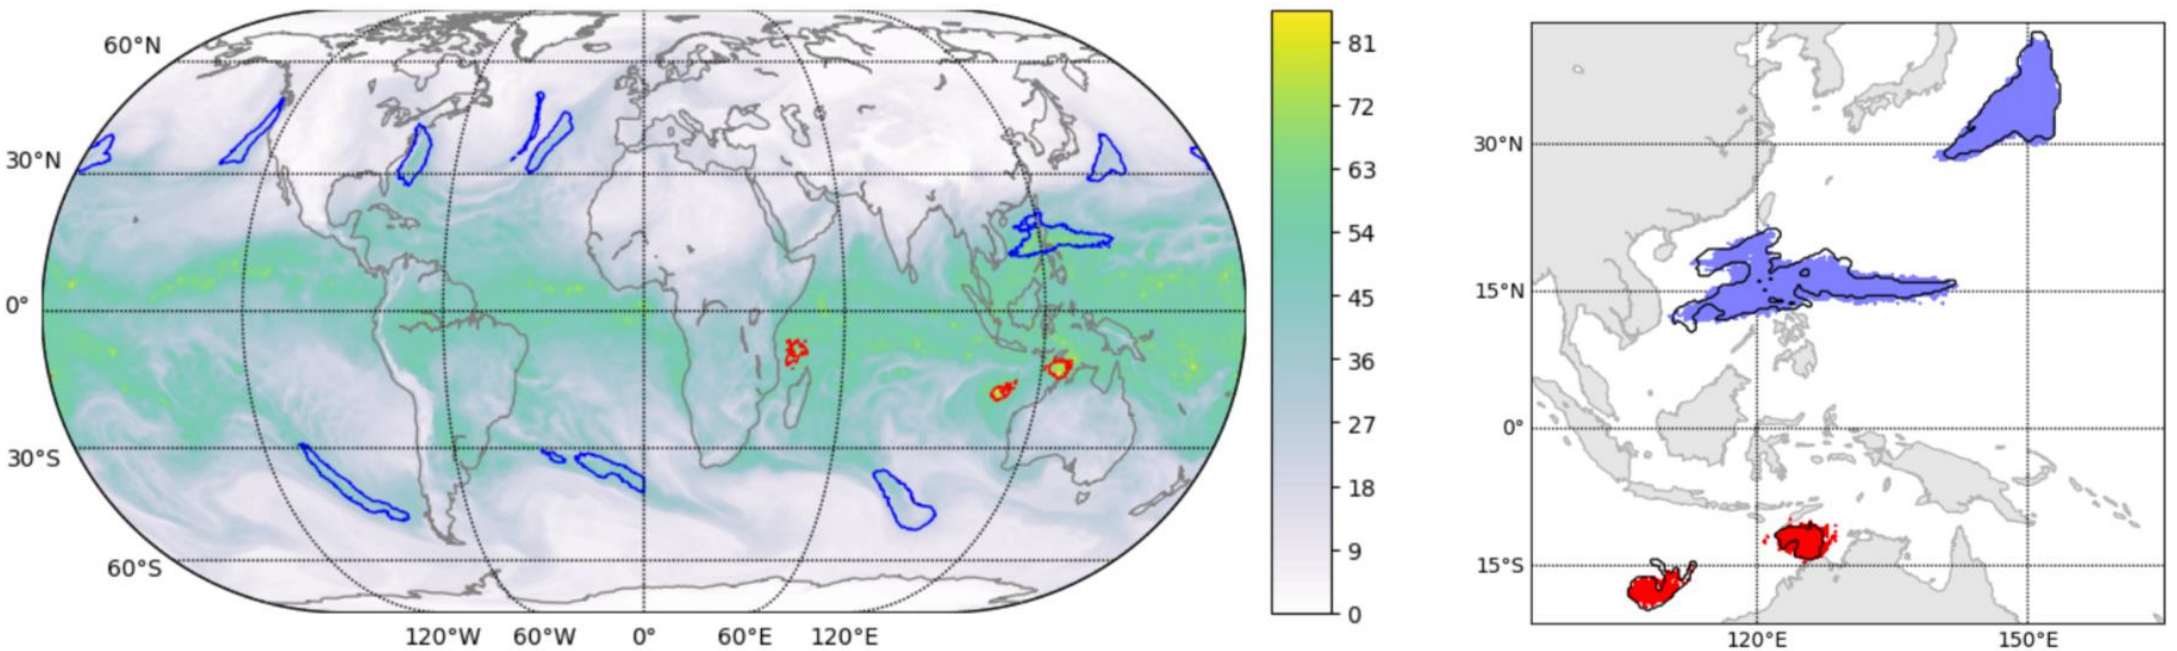

Segmentation of Tropical Storms and Atmospheric Rivers on Summit using convolutional neural networks.

SC18 Gordon Bell Award: NERSC and NVIDIA Team

Exascale Deep Learning for Climate Analytics

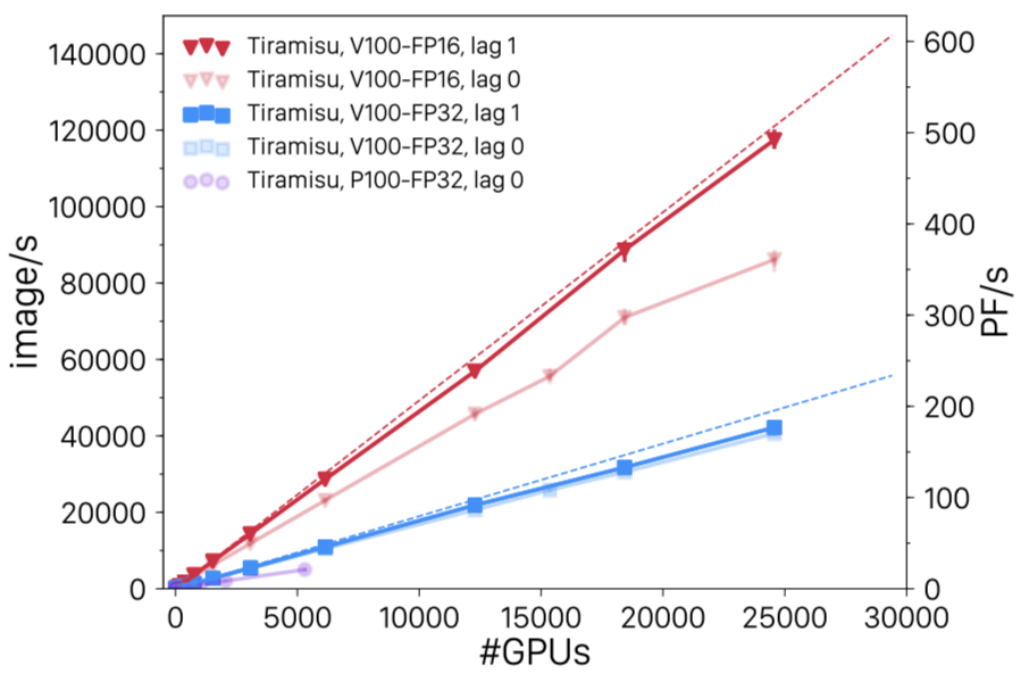

(a) Tiramisu

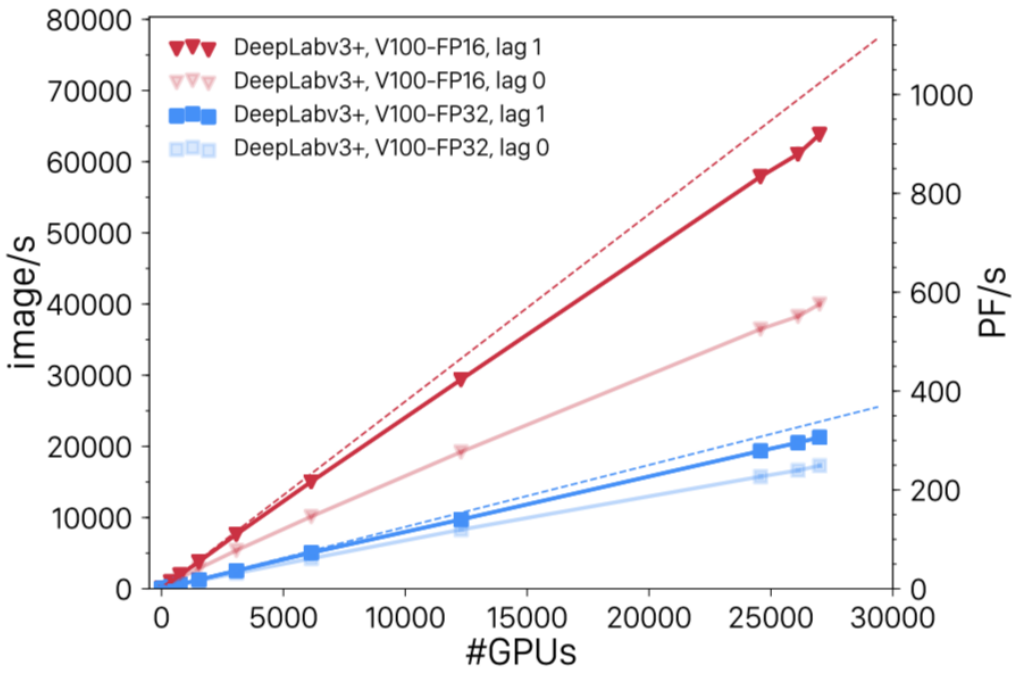

(b) DeepLabv3+

Nearly perfect weak scaling up to 25k GPUS. 1 Exa-flop of performance. 100 years of climate model data in hours
Demonstrates the power of this approach for large-scale data analysis

SC18 NVIDIA Announcements on NWP Models

GPU Accelerated Weather Applications

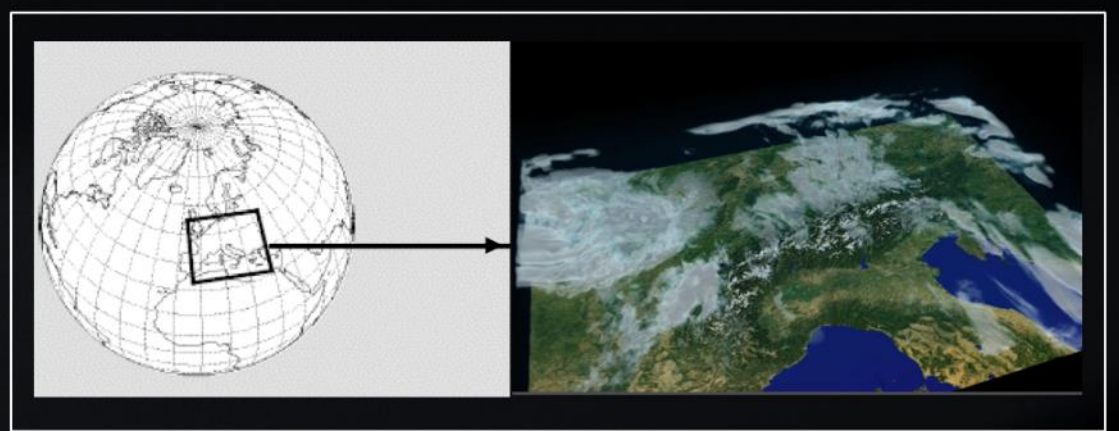

COSMO	WRF	MPAS
14x	9x	12x

* Speedup comparing one 2x Skylake vs 4xV100 Server

SuperComputing 2018

CEO JENSEN HUANG TAKES THE STAGE AT SUPERCOMPUTING 2018

November 12 | Dallas, Texas

*Speedups comparing 2 x Skylake CPU vs. 4 x V100 GPU

NVIDIA

New NVIDIA AI Tech Centre at Reading University

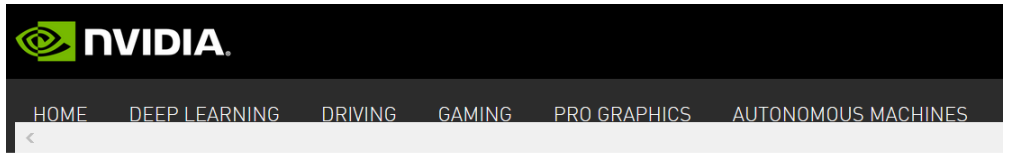

University of Reading

NVIDIA Launches Technology Center to Advance AI Research

ACES

https://blogs.nvidia.com/blog/2019/06/19/ai-technology-center-uk/

The Advanced Computing for Environmental Science (ACES) research group conducts cutting-edge research in computer science to accelerate environmental science.

Environmental science depends on the analysis of large volumes of observational data and on sophisticated simulation schemes, coupling different physics on multiple time and special scales, demanding both supercomputing and specialised data analysis systems. ACES research themes address the future of the relevant computing and data systems.

ACES is based in the Computer Science Department at the University of Reading.

TOPICS

- ACHIEVEMENTS IN HPC AND AI
- **NUMERICAL MODEL DEVELOPMENTS**
- GPU UPDATE ON MPAS-A

NVIDIA Collaborations With Atmospheric Models

	Model	Organizations	Funding Source
Global:			
	E3SM-EAM, SAM	US DOE: ORNL, SNL	E3SM, ECP
	MPAS-A	NCAR, UWyo, KISTI, IBM	WACA II
	FV3/UFS	NOAA	SENA
	NUMA/NEPTUNE	US Naval Res Lab, NPS	ONR
	IFS	ECMWF	ESCAPE
	GungHo/LFRic	MetOffice, STFC	PSyclone
	ICON	DWD, MPI-M, CSCS, MCH	PASC ENIAC
	KIM	KIAPS	KMA
	CLIMA	CLiMA (NASA JPL, MIT, NPS)	Private, US NSF
	FV3	Vulcan, UW/Bretherton	Private
Regional:			
	COSMO	MCH, CSCS, DWD	PASC GridTools
	AceCAST-WRF	TempoQuest	Venture backed

WRFg Collaboration with TempoQuest

- **WRFg Based on ARW release 3.8.1**
- **Several science-ready features:**
  - **Full WRF on GPU; 21 physics options**
  - **Complete nesting functionality**
- **Request download: https://wrfg.net**

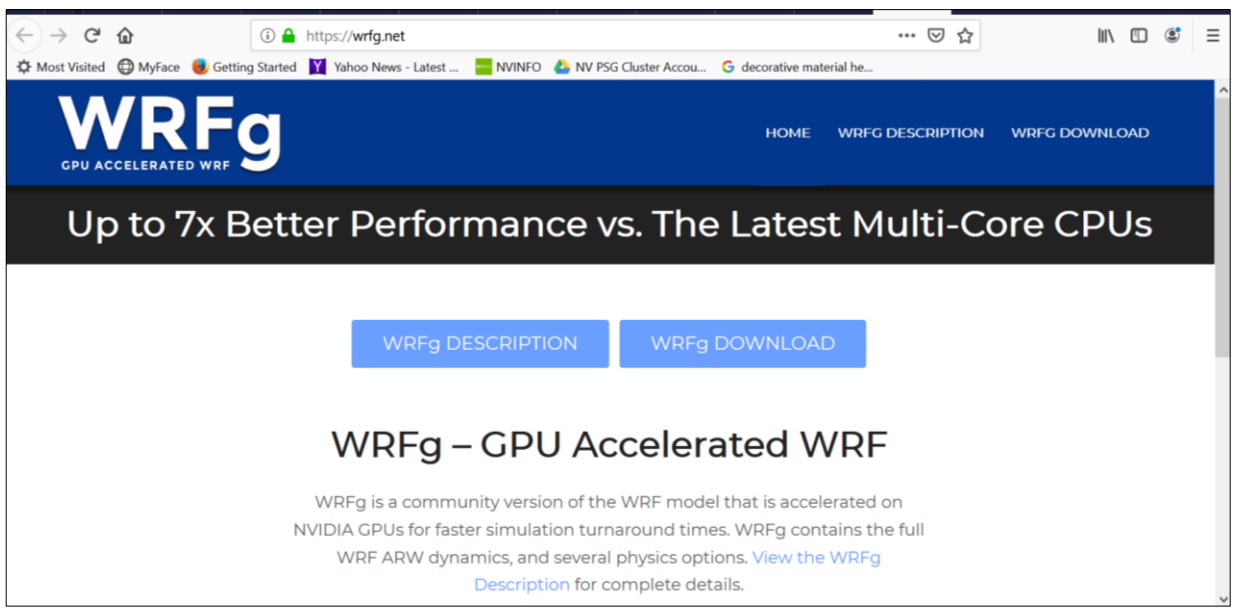

WRFg Physics Options (21)

Microphysics	Option
Kessler	1
WSM6	6
Thompson	8
Morrison	10
Aerosol-aware Thompson	28
Radiation	
Dudhia (sw)	1
RRTMG (lw + sw)	4
Planetary boundary layer	
YSU	1
MYNN	5
Surface layer	
Revised MM5	1
MYNN	5
Land surface	
5-layer TDS	1
Unified Noah	2
RUC	3
Cumulus	
Kain-Fritsch	1; 11; 99
BMJ	2
Grell-Deveni	93
GRIMS Shallow Cumulus	SHCU=3

NV-WRFg Summit Scaling on 512 Nodes / 3,072 GPUs

GPU Performance Study for the WRF Model on Summit

-Jeff Adie, NVIDIA, Gökhan Sever, Rajeev Jain, DOE Argonne NL, and Stan Posey, NVIDIA

Multiscale Coupled Urban Systems – PI, C. Catlett

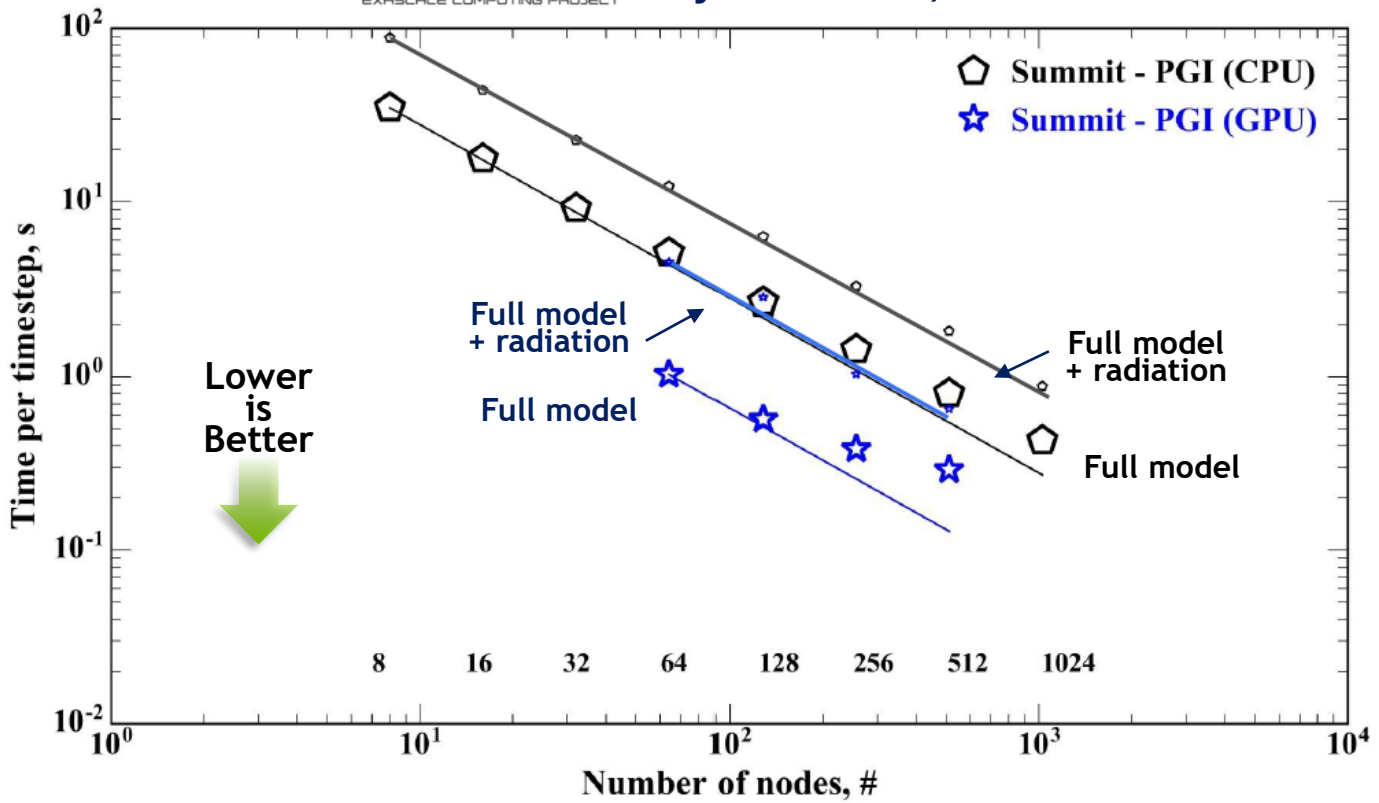

http://www2.mmm.ucar.edu/wrf/users/workshops/WS2019/oral_presentations/3.3.pdf

Joint WRF and MPAS Users' Workshop 2019
NCAR, Boulder, USA

- ORNL Summit node:
  - 2 x P9 + 6 x V100
  - OpenACC, PGI 19.1
- Based on NCAR WRF 3.7.1
- WRF model configuration:
  - Total 3.7B cells
  - Thompson MP
  - RRTM / Dudhia
  - YSU PBL
  - Revised MM5+TDS4

MeteoSwiss Operational COSMO NWP on GPUs

FRANKFURT, Germany, June 18, 2019 (GLOBE NEWSWIRE) -

Swiss Federal Office of Meteorology and Climatology Advances Weather Forecasting With New Cray Supercomputer and Storage

CSCS' new CS-Storm® is configured with 18 compute nodes, each with 8 NVIDIA® V100 GPUs and 2 Intel® Xeon® Gold 6134 CPUs, and includes two Cray ClusterStor® L300 storage systems. The CS-Storm was accepted in April 2019 and will become fully operational in 2020.

18 Nodes x 8 x V100 = 144 Total GPUs

MeteoSwiss Roadmap

- New V100 system in 2019
- New EPS configurations operational in 2020
- New ICON-LAM in ~2022 (Pre-operational in 2020)

MeteoSwiss COSMO NWP Configurations During 2020

With V100 GPUs

IFS from ECMWF

4 per day, 18km / 9km (?)

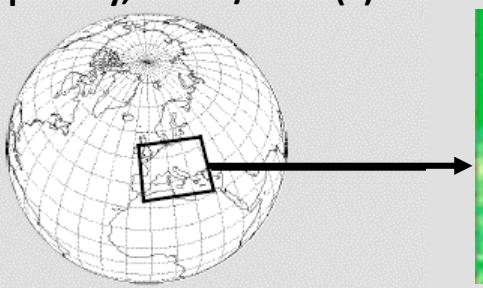

COSMO-2E (2 KM)

4 per day, 5 day forecast

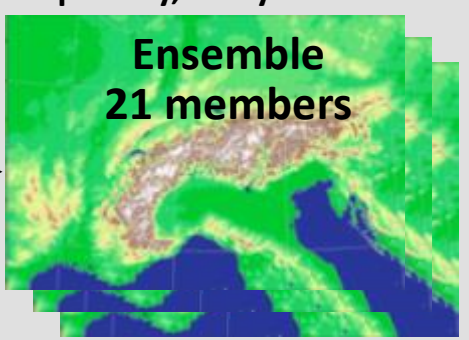

COSMO-1E (1 KM)

8 per day, 33 hr forecast

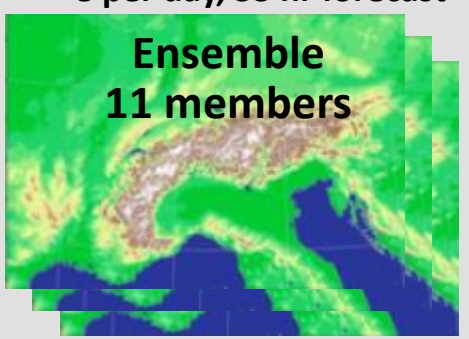

COSMO 1km Near-Global Atmosphere on GPUs

Geosci. Model Dev., 11, 1665–1681, 2018
https://doi.org/10.5194/gmd-11-1665-2018

Geoscientific Model Development
Open Access
EGU

Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0

Oliver Fuhrer[1], Tarun Chadha[2], Torsten Hoefler[3], Grzegorz Kwasniewski[3], Xavier Lapillonne[1], David Leutwyler[4], Daniel Lüthi[4], Carlos Osuna[1], Christoph Schär[4], Thomas C. Schulthess[5,6], and Hannes Vogt[6]

[1]Federal Institute of Meteorology and Climatology, MeteoSwiss, Zurich, Switzerland
[2]ITS Research Informatics, ETH Zurich, Switzerland
[3]Scalable Parallel Computing Lab, ETH Zurich, Switzerland
[4]Institute for Atmospheric and Climate Science, ETH Zurich, Switzerland
[5]Institute for Theoretical Physics, ETH Zurich, Switzerland
[6]Swiss National Supercomputing Centre, CSCS, Lugano, Switzerland

Correspondence: Oliver Fuhrer (oliver.fuhrer@meteoswiss.ch)

Received: 16 September 2017 – Discussion started: 5 October 2017
Revised: 7 February 2018 – Accepted: 8 February 2018 – Published: 2 May 2018

Abstract. The best hope for reducing long-standing global climate model biases is by increasing resolution to the kilometer scale. Here we present results from an ultrahigh-resolution non-hydrostatic climate model for a near-global setup running on the full Piz Daint supercomputer on 4888 GPUs (graphics processing units). The dynamical core of the model has been completely rewritten using a domain-specific language (DSL) for performance portability across

in the availability of water resources and the occurrence of droughts (Pachauri and Meyer, 2014).

Current climate projections are mostly based on global climate models (GCMs). These models represent the coupled atmosphere–ocean–land system and integrate the governing equations, for instance, for a set of prescribed emissions scenarios. Despite significant progress during the last decades, uncertainties are still large. For example, current estimates

Large Scale COSMO HPC Demonstration Using ~5000 GPUs

MeteoSwiss

ETH zürich

CSCS

Source: https://www.geosci-model-dev-discuss.net/gmd-2017-230/

NVIDIA

COSMO 1km Near-Global Atmosphere on GPUs

Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4'888 GPUs with COSMO 5.0

Oliver Fuhrer[1], Tarun Chadha[2], Torsten Hoefler[3], Grzegorz Kwasniewski[3], Xavier Lapillonne[1], David Leutwyler[4], Daniel Lüthi[4], Carlos Osuna[1], Christoph Schär[4], Thomas C. Schulthess[5,6], and Hannes Vogt[6]

- Oliver Fuhrer, et al, MeteoSwiss

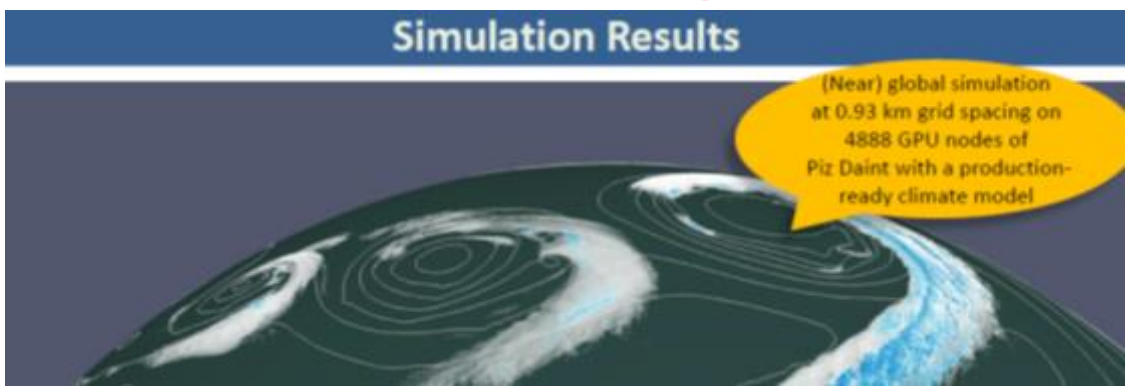

Piz Daint
#6 Top500
25.3 PetaFLOPS
5320 x P100 GPUs

Strong Scaling to 4,888 x P100 GPUs

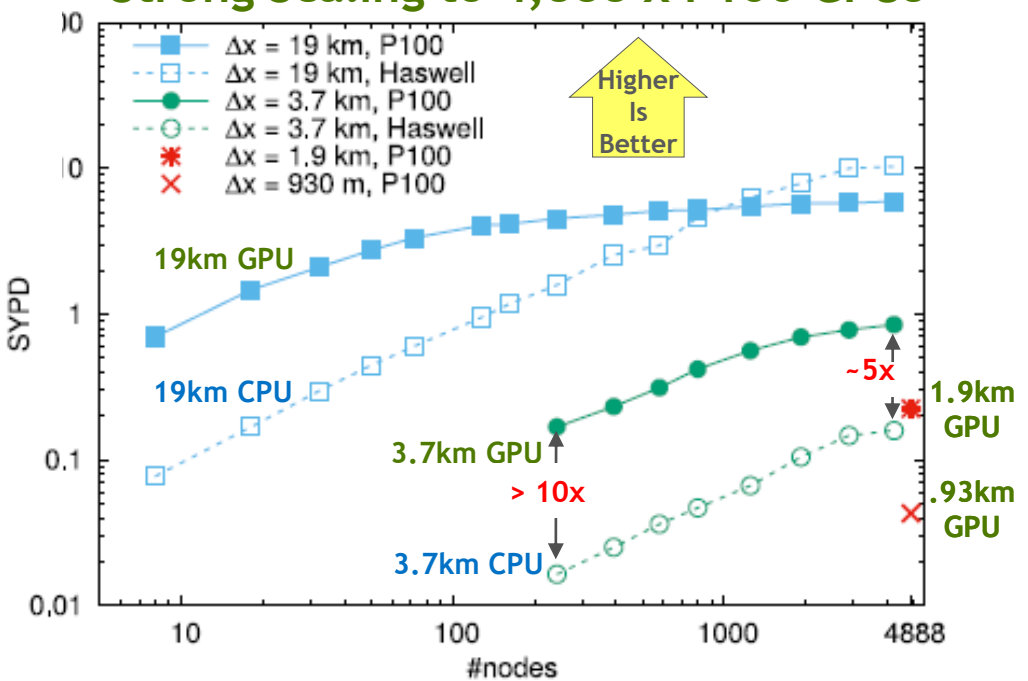

Δx	# nodes	SYPD	MWh / SY
0.93 km	4,888	0.043	596
1.9 km	4,888	0.23	97.8

Source: https://www.geosci-model-dev-discuss.net/gmd-2017-230/

NOAA FV3 and GPU Developments (Chen - 2019)

NOAA FV3 GPU Strategy Includes OpenACC and GridTools

-From Presentation by Dr. Xi Chen, NOAA GFDL, PASC 19, June 2019, Zurich, CH

FV3 GPU development status and external collaborations

- OpenACC + Cuda
  - In house **[X. Chen, others]**
- GridTools/Kokkos
  - Swiss National Supercomputing Centre (CSCS)
  - Vulcan Group **[C. Bretherton, O. Fuhrer]**
  - DOE/NASA **[W. Putman, others]**

- **2012: Early GPU development by NASA GSFC GMAO**

PASC19 Zurich Switzerland 12-14 June 2019

Future FV3 Developments and the Participation of DYAMOND in the Era of E-Class HPCs

Xi Chen[1,2], Shian-Jiann Lin[2], and the FV3 Team

13 June 2019

Zürich, Switzerland

[1]Atmospheric Oceanic Sciences, Princeton University, Princeton, NJ USA
[2]NOAA's Geophysical Fluid Dynamics Lab, Princeton, NJ USA

2019 ORNL Hackathons and GPU Model Progress

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

ABOUT OLCF ▾ OLCF RESOURCES ▾ R&D ACTIVITIES ▾ SCIENCE AT OLCF ▾ FOR USERS ▾ OLCF MEDIA ▾

Location - Date	Organizations	Model(s)	Hackathon Project
KISTI (KR) - Feb	KISTI	MPAS	Physics (WSM6)
CAS (CN) - May	CMA	GRAPES	PRM advection
ETH Zurich- Jun	MCH, MPI-M, CSCS	ICON	Physics, radiation
MIT - Jun	MIT, CliMA	CliMA Ocean	Subgrid scale LES
Princeton - Jun	NOAA GFDL	FV3/UFS	SWE mini-app kernels, UFS radiation package
NERSC - Jul	DOE LBNL	E3SM	MMF (ECP)
Met Office - Sep	Met Office, STFC	NEMOVAR, WW III	Miniapp (?)
ORNL - Oct	DOE ORNL, SNL	E3SM	SCREAM (Kokkos)

https://www.olcf.ornl.gov/for-users/training/gpu-hackathons/

2019 ORNL Hackathons and GPU Model Progress

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

ABOUT OLCF ▾ OLCF RESOURCES ▾ R&D ACTIVITIES ▾ SCIENCE AT OLCF ▾ FOR USERS ▾ OLCF MEDIA ▾

Location - Date	Organizations	Model(s)	Hackathon Project
KISTI (KR) - Feb	KISTI	MPAS	Physics (WSM6)
CAS (CN) - May	CMA	GRAPES	PRM advection
ETH Zurich- Jun	MCH, MPI-M, CSCS	ICON	Physics, radiation
MIT - Jun	MIT, CliMA	CliMA Ocean	Subgrid scale LES
Princeton - Jun	NOAA GFDL	FV3/UFS	SWE mini-app kernels, UFS radiation package
NERSC - Jul	DOE LBNL	E3SM	MMF (ECP)
Met Office - Sep	Met Office, STFC	NEMOVAR, WW III	Miniapp (?)
ORNL - Oct	DOE ORNL, SNL	E3SM	SCREAM (Kokkos)

https://www.olcf.ornl.gov/for-users/training/gpu-hackathons/

CliMA: New Climate Model Development

CliMA
CLIMATE MODELING ALLIANCE

https://clima.caltech.edu/

NVIDIA
DRIVERS • PRODUCTS • DEEP LEARNING AND AI • COMMUNITIES • SUPPORT SHOP ABOUT NVIDIA •

https://blogs.nvidia.com/blog/2019/07/17/clima-climate-model/

Spotting Clouds on the Horizon: AI Resolves Uncertainties in Climate Projections

July 17, 2019 by ISHA SALIAN

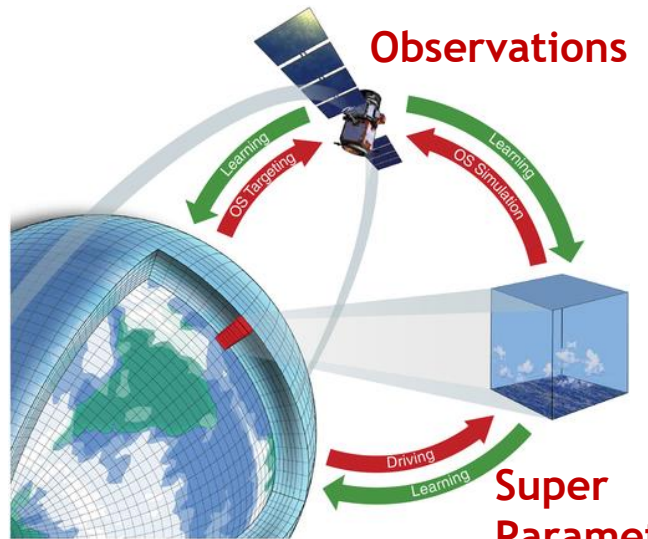

Caltech

JPL
Jet Propulsion Laboratory
California Institute of Technology

Ocean

NPS
PRAESTANTIA PER SCIENTIAM

Atmosphere

NVIDIA

CliMA: New Climate Model Development

Pushing the Envelope in Ocean Modeling

https://clima.caltech.edu/

-From Keynote Presentation by Dr. John Marshall, MIT, at Oxford AI Workshop, Sep 2019, Oxford, UK

CliMA CLIMATE MODELING ALLIANCE

Oceananigans (LES Ocean model)

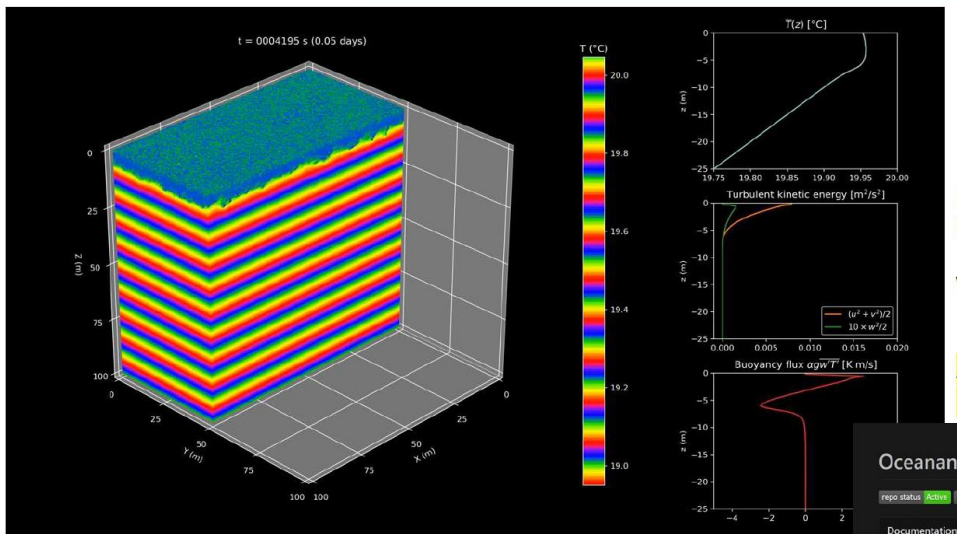

Upper ocean mixing driven by wind and cooling

Finite volume, non-hydrostatic model
LES Anisotropic Minimum Dissipation closure

Written in Julia and supports CPU and

Ali Ramadhan, Chris Hill, J-M Car
Greg Wagner & Valentin Churavy

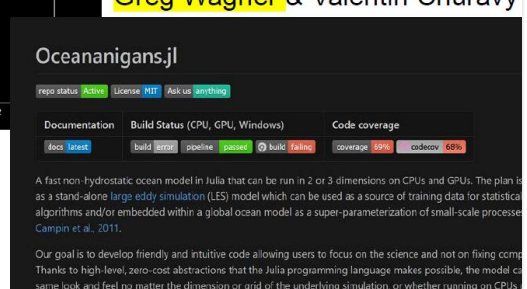

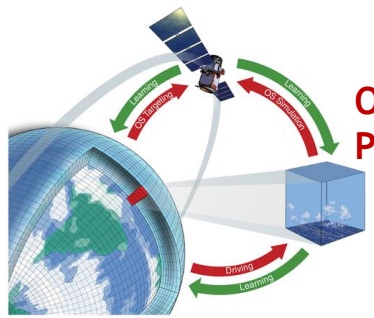

Ocean LES Super Parameterization

CliMA CLIMATE MODELING ALLIANCE

Pushing the Envelope in Ocean Modeling

MIT Ocean Modeling Group
Raffaele Ferrari and John Marshall

Brandon Allen, Jean-Michel Campin, Basile Gallet, Chris Hill, Ali Ramadhan, Andre Souza, Greg Wagner.

MIT JuliaLab
Alan Edelman
Valentin Churavy

Developing Ocean component of CliMA

Developing an Earth System model that automatically learns from data and targeted high-resolution simulations (of. e.g. clouds and ocean mixing)

Tapio Schneider, Andrew Stewart
CalTech

SCHMIDT FUTURES · EARTHRISE · PAUL G. ALLEN PHILANTHROPIES · HEISING-SIMONS FOUNDATION

OpenACC GPU Development for LFRic Model

- **OpenACC collaboration with MetOffice and SFTC: LFRic model**
  - GungHo-MV (matrix-vector operations) OpenACC kernel developed by MetOffice
  - NVIDIA optimizations applied to the OpenACC kernel achieved 30x improvement
  - Improved OpenACC code provided to STFC as 'target' for Psyclone auto-generation

Met Office

"Optimization of an OpenACC Weather Simulation Kernel"

- A. Gray, NVIDIA

30x Improvement from NVIDIA optimizations over original MetO code

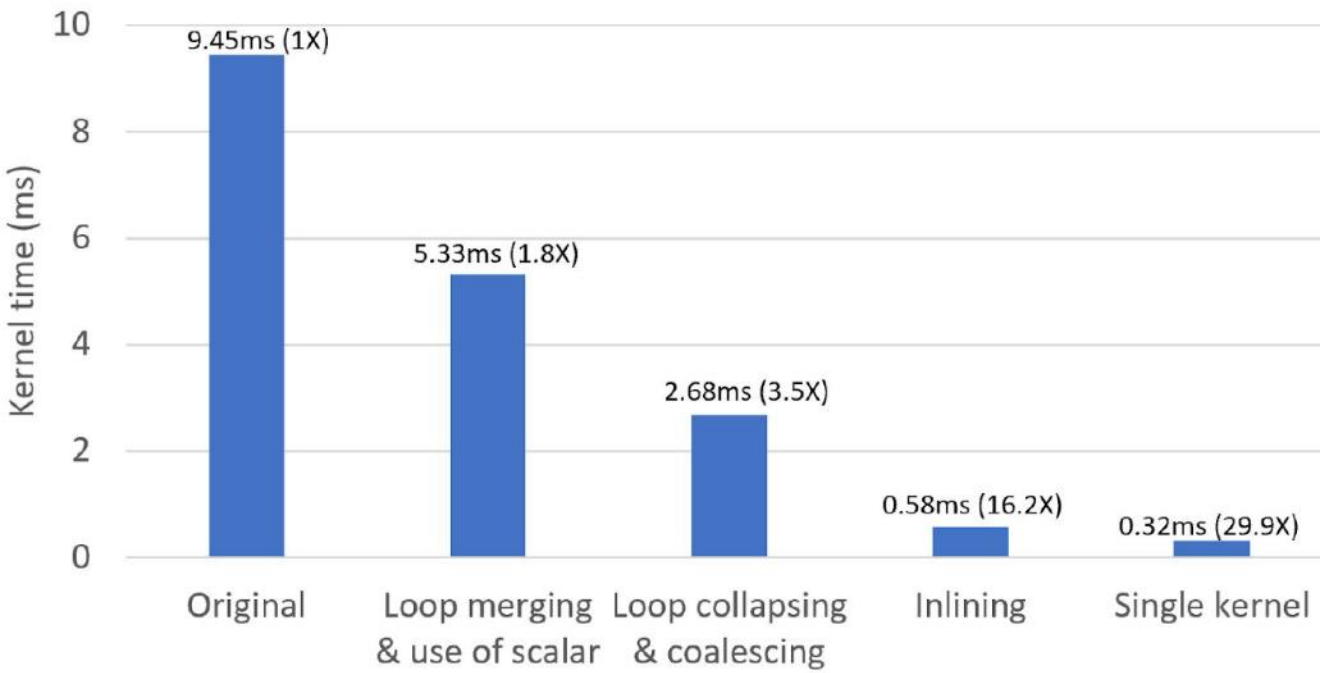

https://www.openacc.org/blog/optimizing-openacc-weather-simulation-kernel

OpenACC GPU Development for LFRic Model

- **OpenACC collaboration with MetOffice and SFTC: LFRic model**
  - GungHo-MV (matrix-vector operations) OpenACC kernel developed by MetOffice
  - NVIDIA optimizations applied to the OpenACC kernel achieved 30x improvement
  - Improved OpenACC code provided to STFC as 'target' for Psyclone auto-generation

"Optimization of an OpenACC Weather Simulation Kernel"

- A. Gray, NVIDIA

30x Improvement from NVIDIA optimizations over original MetO code

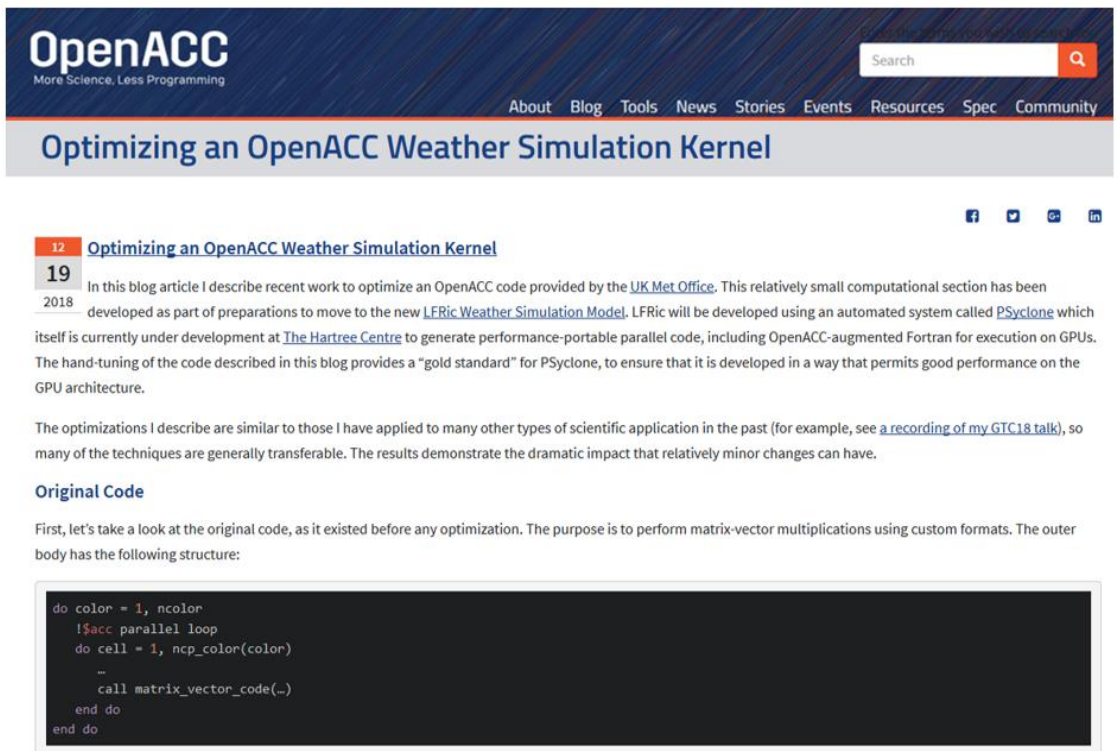

https://www.openacc.org/blog/optimizing-openacc-weather-simulation-kernel

ESCAPE Development of Weather & Climate Dwarfs

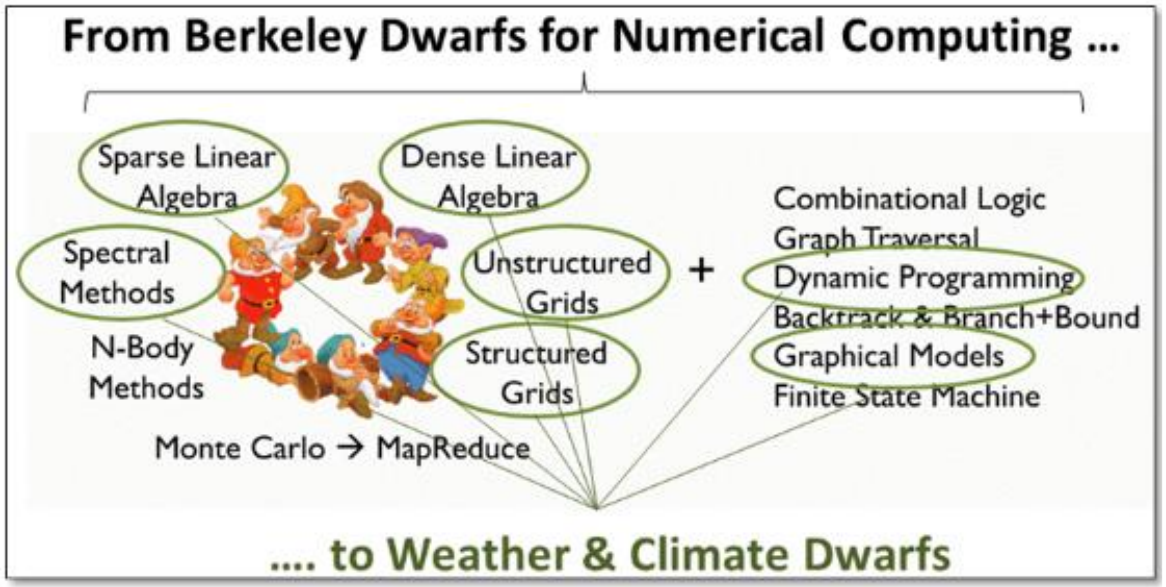

ICHEC · Deutscher Wetterdienst · Optalysys · Bull · PSNC · Loughborough University · MeteoSwiss · dmi · RMI · NVIDIA · METEO FRANCE Toujours un temps d'avance

NVIDIA-Developed Dwarf: Spectral Transform - Spherical Harmonics

Batched 1D FFT variable length

Batched Legendre Transform (GEMM) variable length

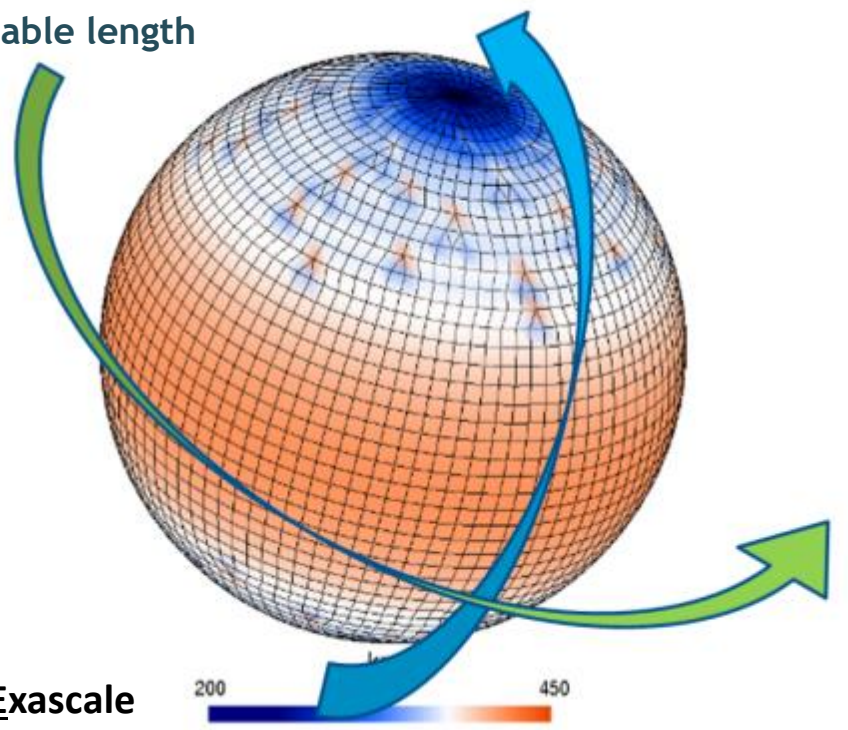

ESCAPE = Energy efficient SCalable Algorithms for weather Prediction at Exascale

ECMWF IFS Dwarf Optimizations – Single-GPU

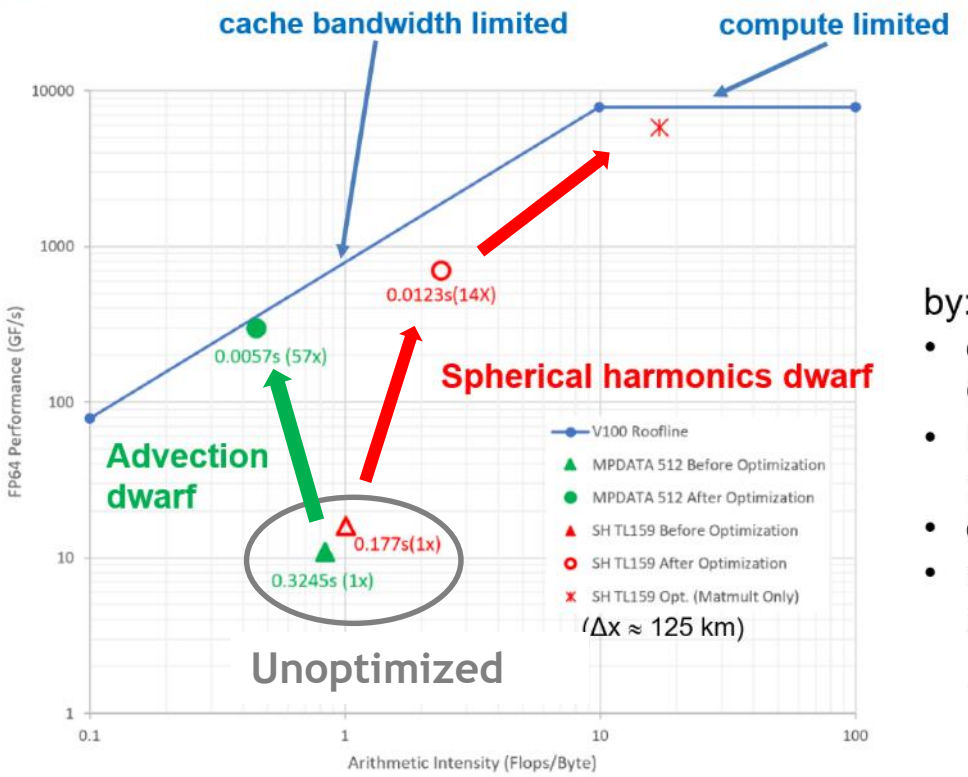

SH Dwarf = 14x

Advection Dwarf = 57x

by:

- exposing parallelism in loops for OpenACC mapping
- Kernel optimization by memory mapping
- exploiting CUDA BLAS features
- minimizing data allocation and movement

From *"ECMWF Scalability Programme"*

Dr. Peter Bauer, UM User Workshop, MetOffice, Exeter, UK 15 June 2018

- **Single V100 GPU improved SH dwarf by 14x vs. original**
- **Single V100 GPU improved MPDATA dwarf 57x vs. orig**

ECMWF IFS SH Dwarf Optimization – Multi-GPU

Hybrid Computing – multiple GPU

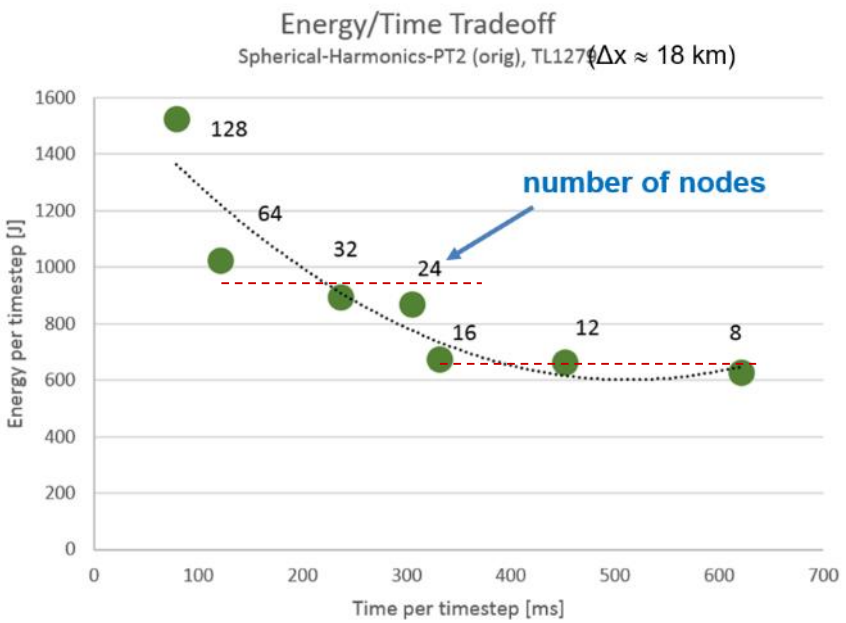

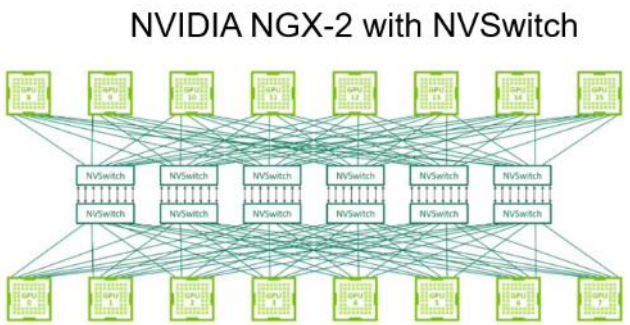

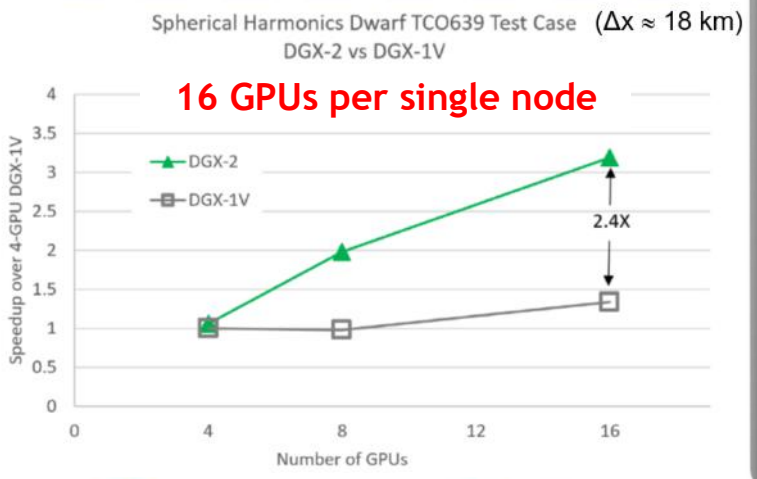

From *"ECMWF Scalability Programme"*

Dr. Peter Bauer, UM User Workshop, MetOffice, Exeter, UK 15 June 2018

- Results of Spherical Harmonics Dwarf on NVIDIA DGX System
- Additional 2.4x gain from DGX-2 NVSwitch for 16 GPU systems

TOPICS

- ACHIEVEMENTS IN HPC AND AI
- NUMERICAL MODEL DEVELOPMENTS
- **GPU UPDATE ON MPAS-A**