


A baseline for global weather and climate simulations at 1 km resolution

Nils P. Wedi

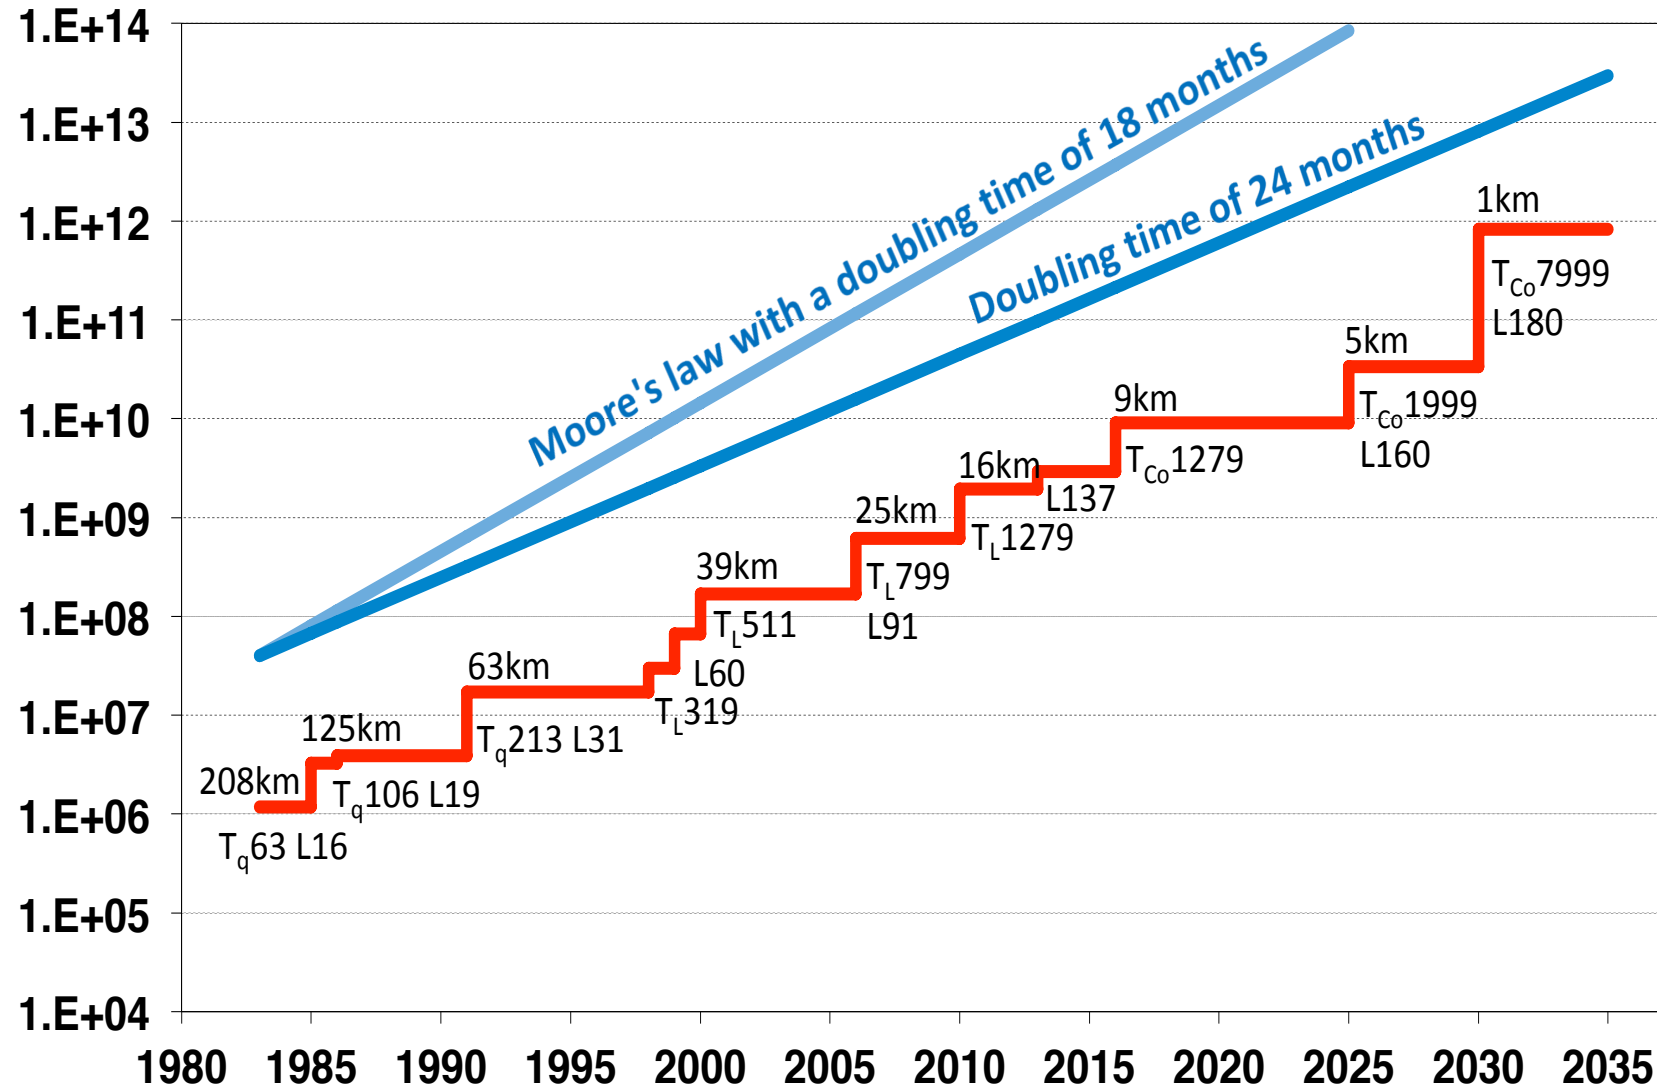
European Centre for Medium-Range Weather Forecasts (ECMWF)



Outline

- Performance, Portability and Complexity evolution of the Integrated Forecasting System (IFS)
- Towards increasing realism with storm-scale resolving simulations
- A first digital twin prototype: a global seasonal simulation with 1.4 km grid-spacing on Summit


Computational power drives spatial resolution



(Schulthess et al, 2019)

ECMWF's progress in degrees of freedom
(levels x grid columns x prognostic variables)

Spectral transform based model at global average 1.4 km grid spacing



6,599,680 points x 137 levels x 10 vars
at ~9km ~ **9 Billion points or ~100TB/day**



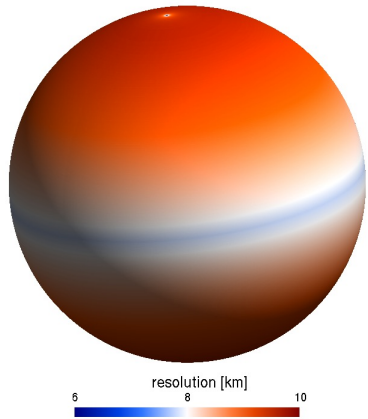
TCo7999 L137
256,800,000 points x 137 levels x 10 variables
at ~1,4km
**352 billion points x 960 pp steps ==
~100TB/simulated month
Summit SIMULATION**

Model aspect	currently operational		
	IFS-FVM	IFS-ST	IFS-ST (NH option)
Equation system	fully compressible	hydrostatic primitive	fully compressible
Prognostic variables	$\rho_d, u, v, w, \theta', \varphi', r_v, r_l, r_r, r_i, r_s$	$\ln p_s, u, v, T_v, q_v, q_l, q_r, q_i, q_s$	$\ln \pi_s, u, v, d_4, T_v, \hat{q}, q_v, q_l, q_r, q_i, q_s$
Horizontal coordinates	λ, ϕ (lon–lat)	λ, ϕ (lon–lat)	λ, ϕ (lon–lat)
Vertical coordinate	generalized height	hybrid sigma–pressure	hybrid sigma–pressure
Horizontal discretization	unstructured finite volume (FV)	spectral transform (ST)	spectral transform (ST)
Vertical discretization	structured FD–FV	structured FE	structured FD or FE
Horizontal staggering	co-located	co-located	co-located
Vertical staggering	co-located	co-located	co-located, Lorenz
Horizontal grid	octahedral Gaussian or arbitrary	octahedral Gaussian	octahedral Gaussian
Time stepping scheme	2-TL SI	2-TL constant-coefficient SI	2-TL constant-coefficient SI with ICI
Advection	conservative FV Eulerian	non-conservative SL	non-conservative SL


The Finite-Volume Module of the IFS (IFS-FVM) provides **complementary features**, e.g. *local computational patterns and predominantly nearest neighbour, non-hydrostatic, complex and steep orography, conservation, mesh adaptation.*

A "non-disruptive" evolutionary approach in dynamical core development

Target: Octahedral grid O8000



Atlas mesh generation and flexibility through alternative function spaces



The next 10 years ...

Non-intrusive introduction of Atlas library in IFS

- Flexible data structures to facilitate future developments
- Capability to link CPU to GPU host
- Mathematical operators, grid generation & multiple grids capability

Code modularity to enable seamless use of heterogenous architectures (CPU/GPU)

Continue improving IFS-ST core algorithms

- Cubic octahedral grid (efficient and more scalable)
- Fast LT algorithm
- Continue improving SL advection & its TL/AD

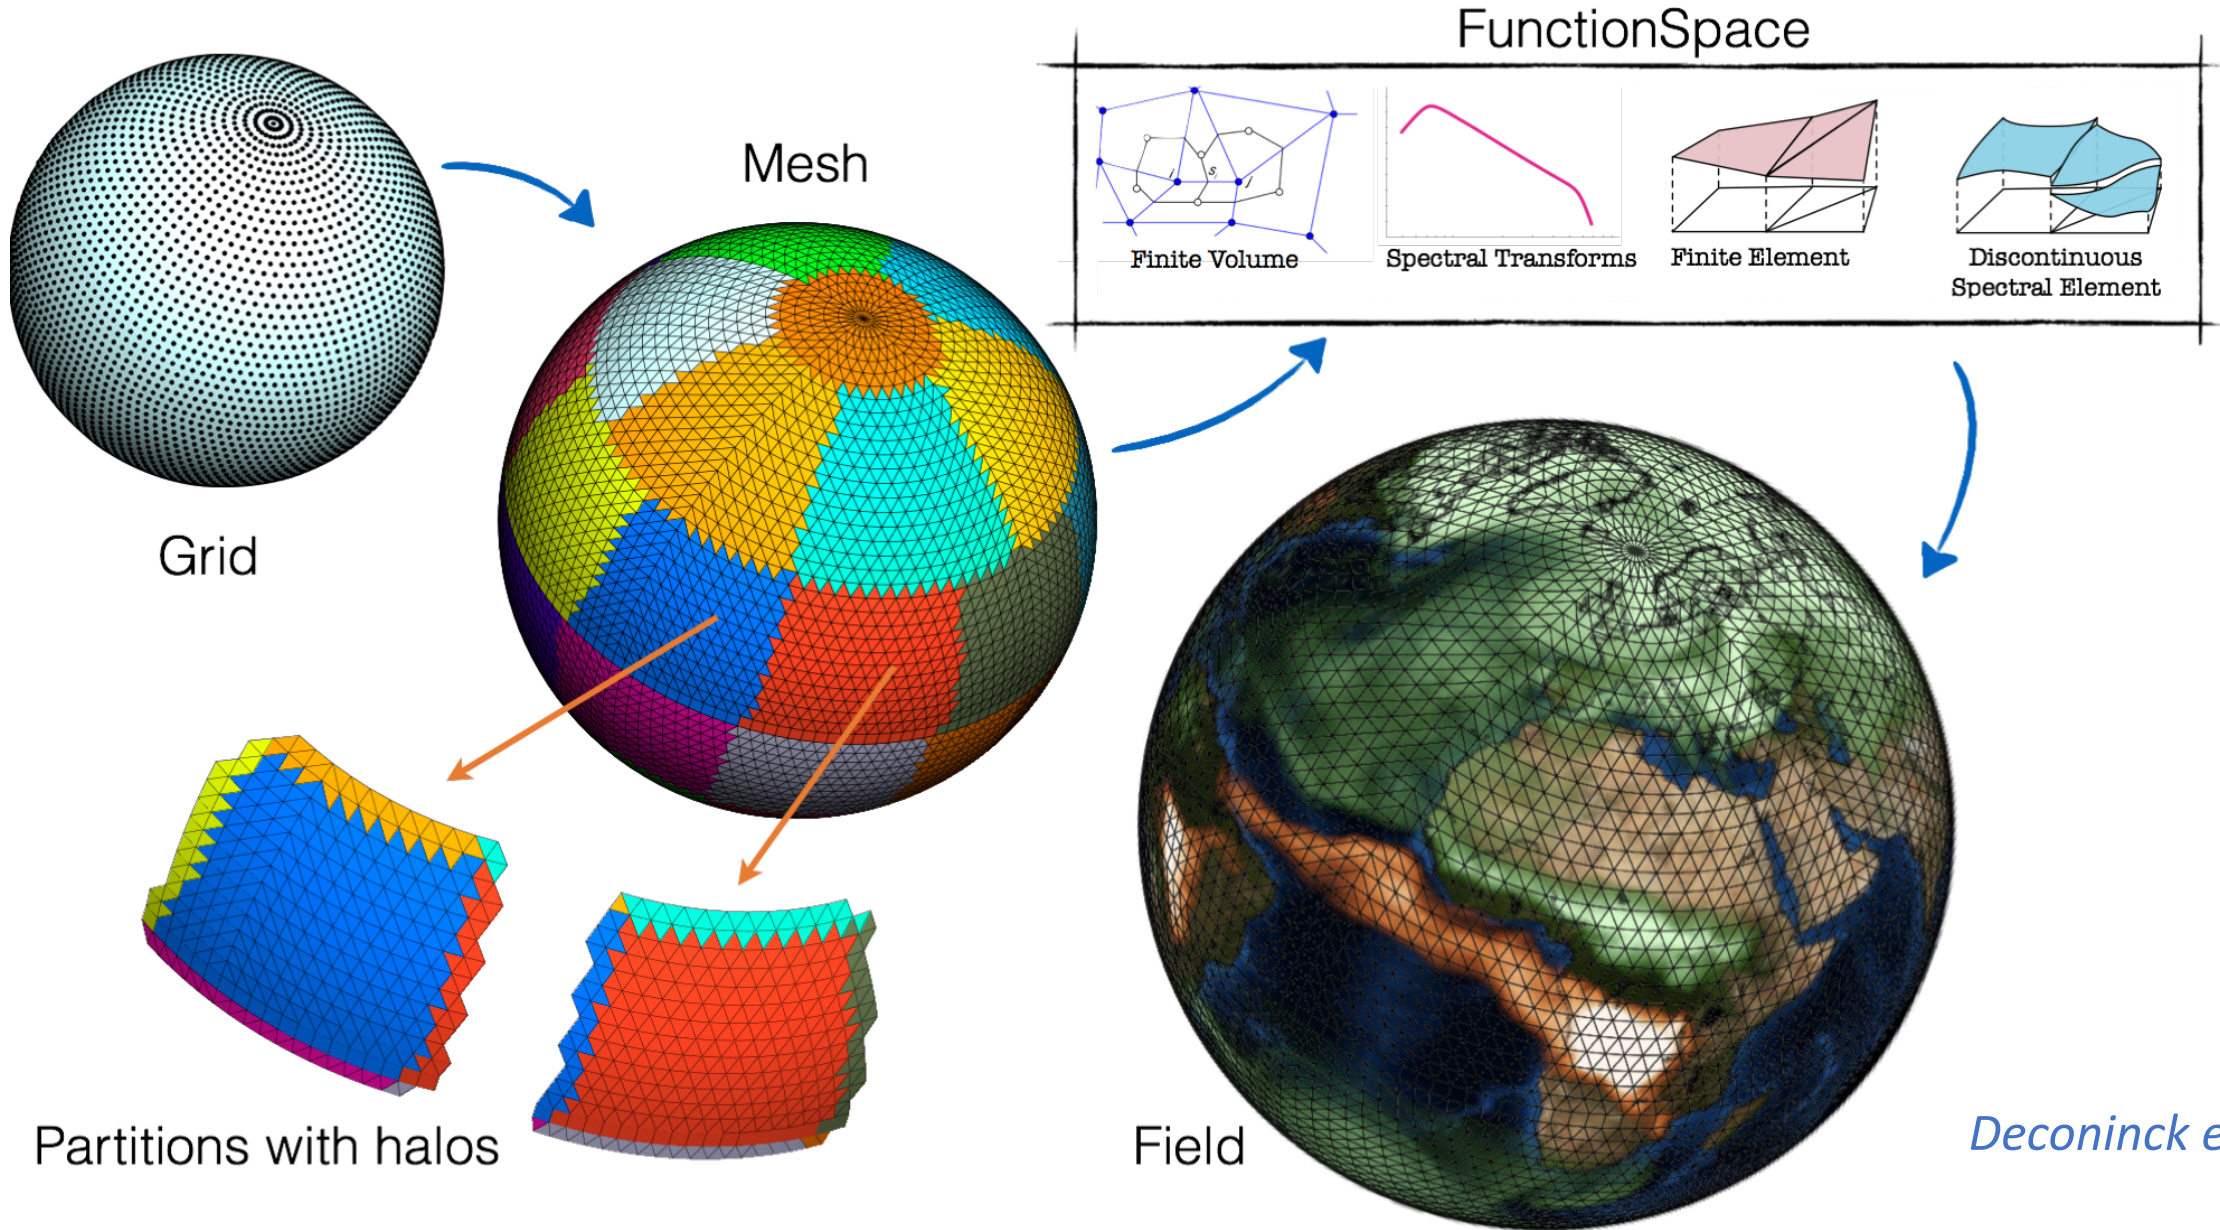
Developing NH FVM on same O-grid advantageous for finest resolutions:

- Low communication (compact stencil) and scalability
- Steep slope Mass conservation and good accuracy

As resolution keeps increasing the "triple crossing point" is reached:

- NH formulation is advantageous
- FVM more efficient and more adaptive via DSL than IFS-ST
- FVM delivers the same quality of forecasts

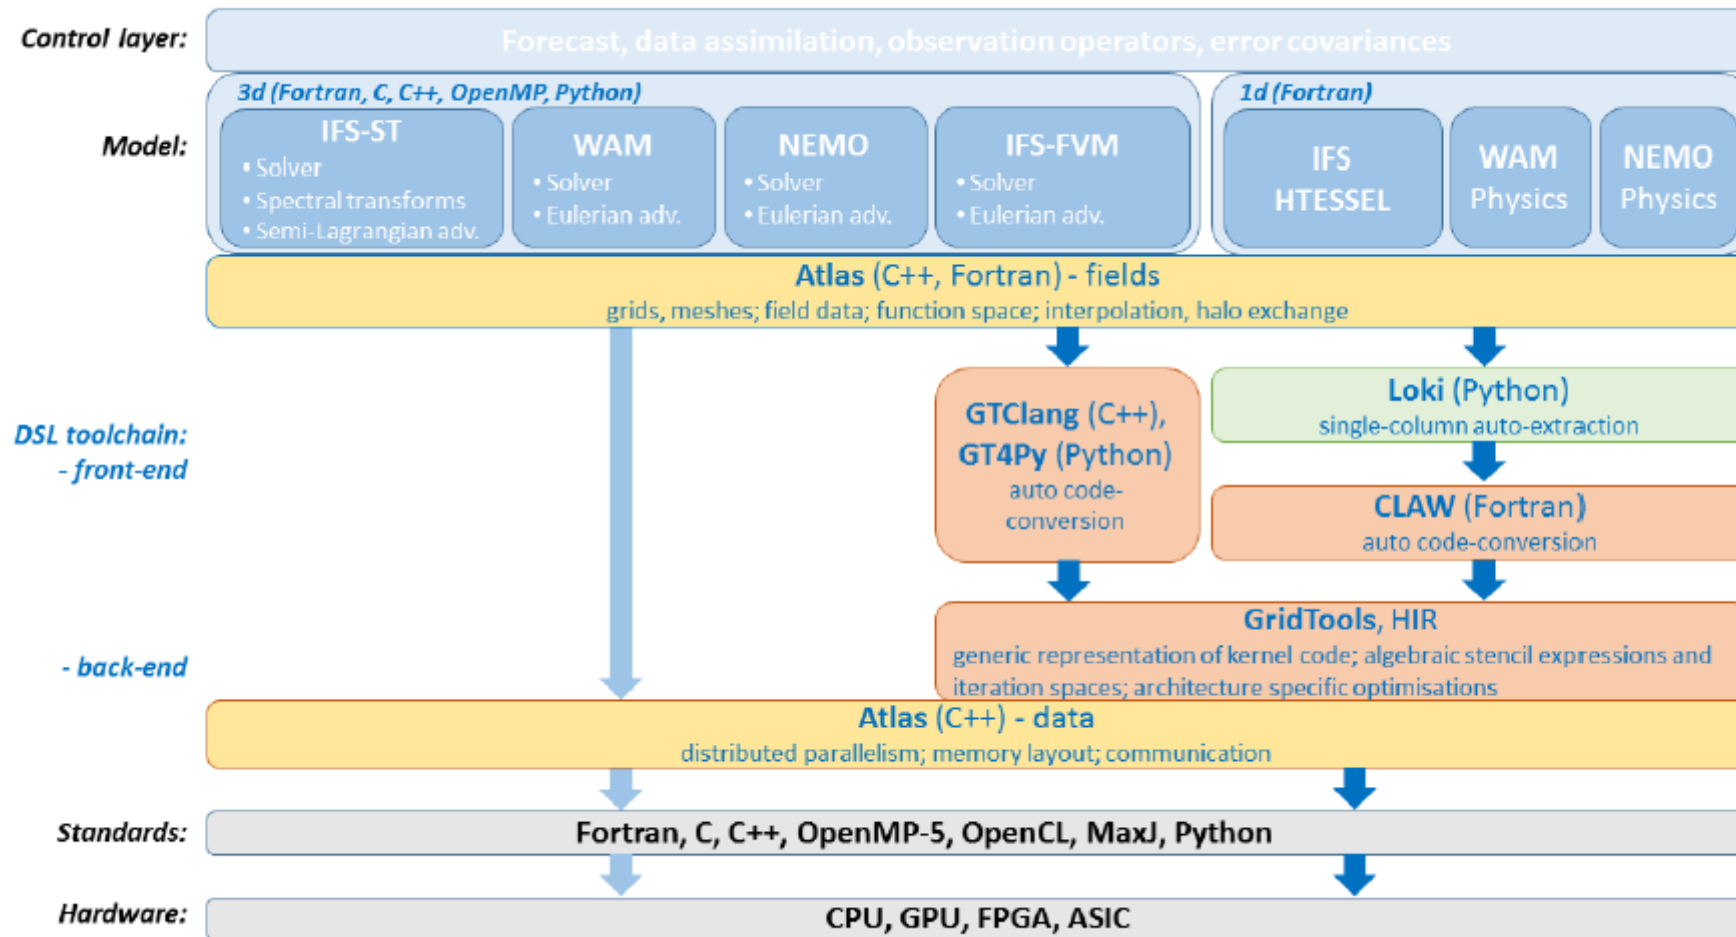
FVM operational



Deconinck et al. 2017

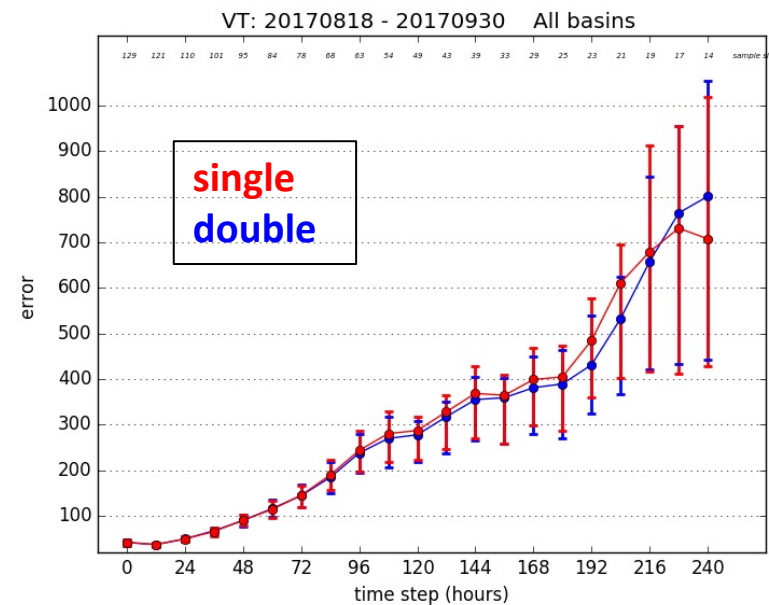
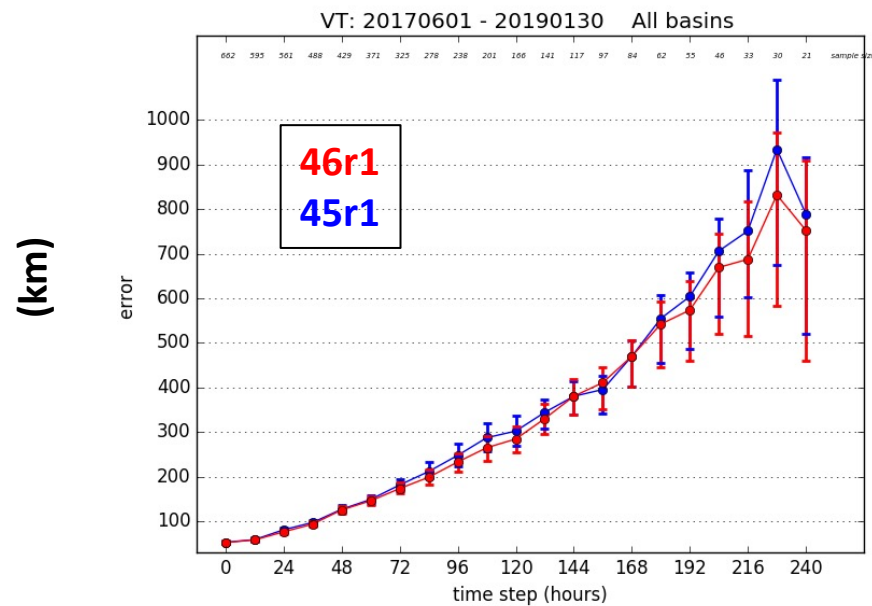
Performance and portability

M. Lange, O. Marsden, B. Reuter

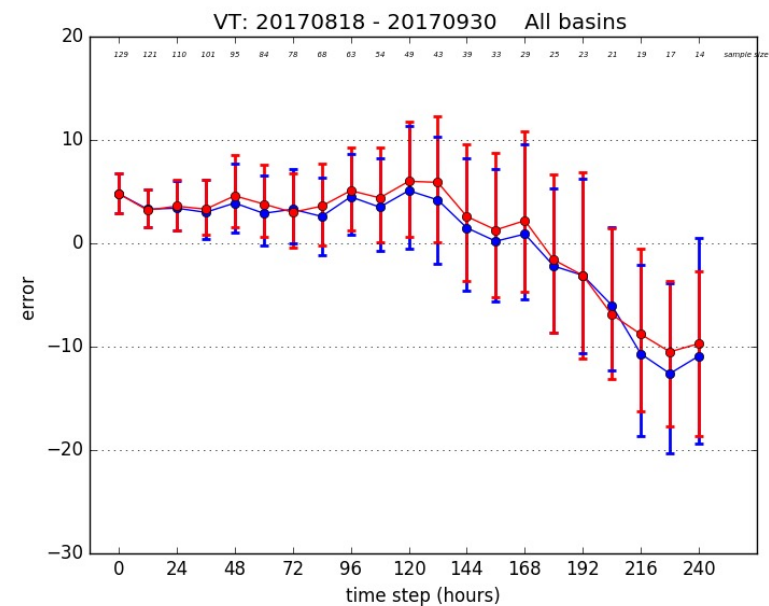
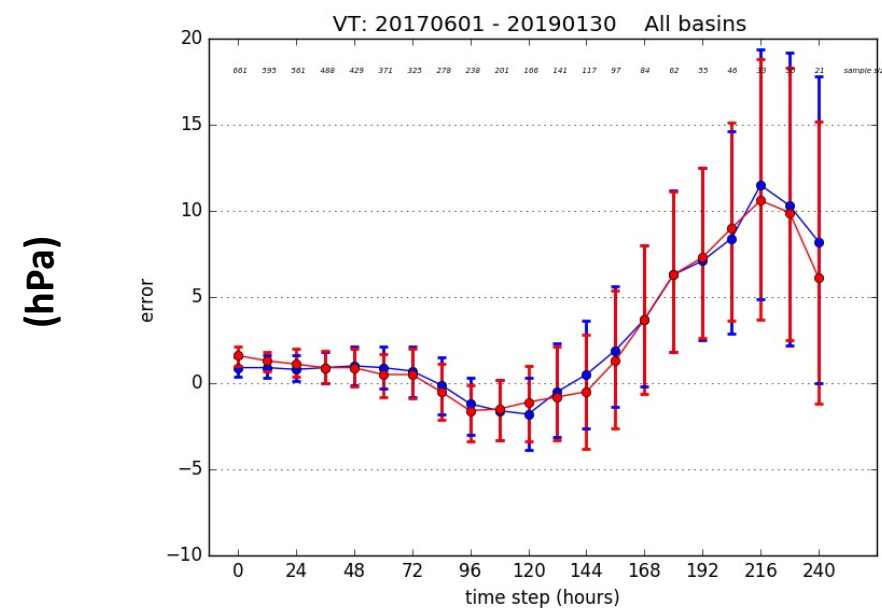


Structure and components necessary for the transition of the IFS to separate applied science from hardware sensitive code level

Single versus double precision tropical Cyclone error measures




mean position error



mean intensity error

ECMWF Scalability Programme



DIGITAL TWINS

Technical Memo



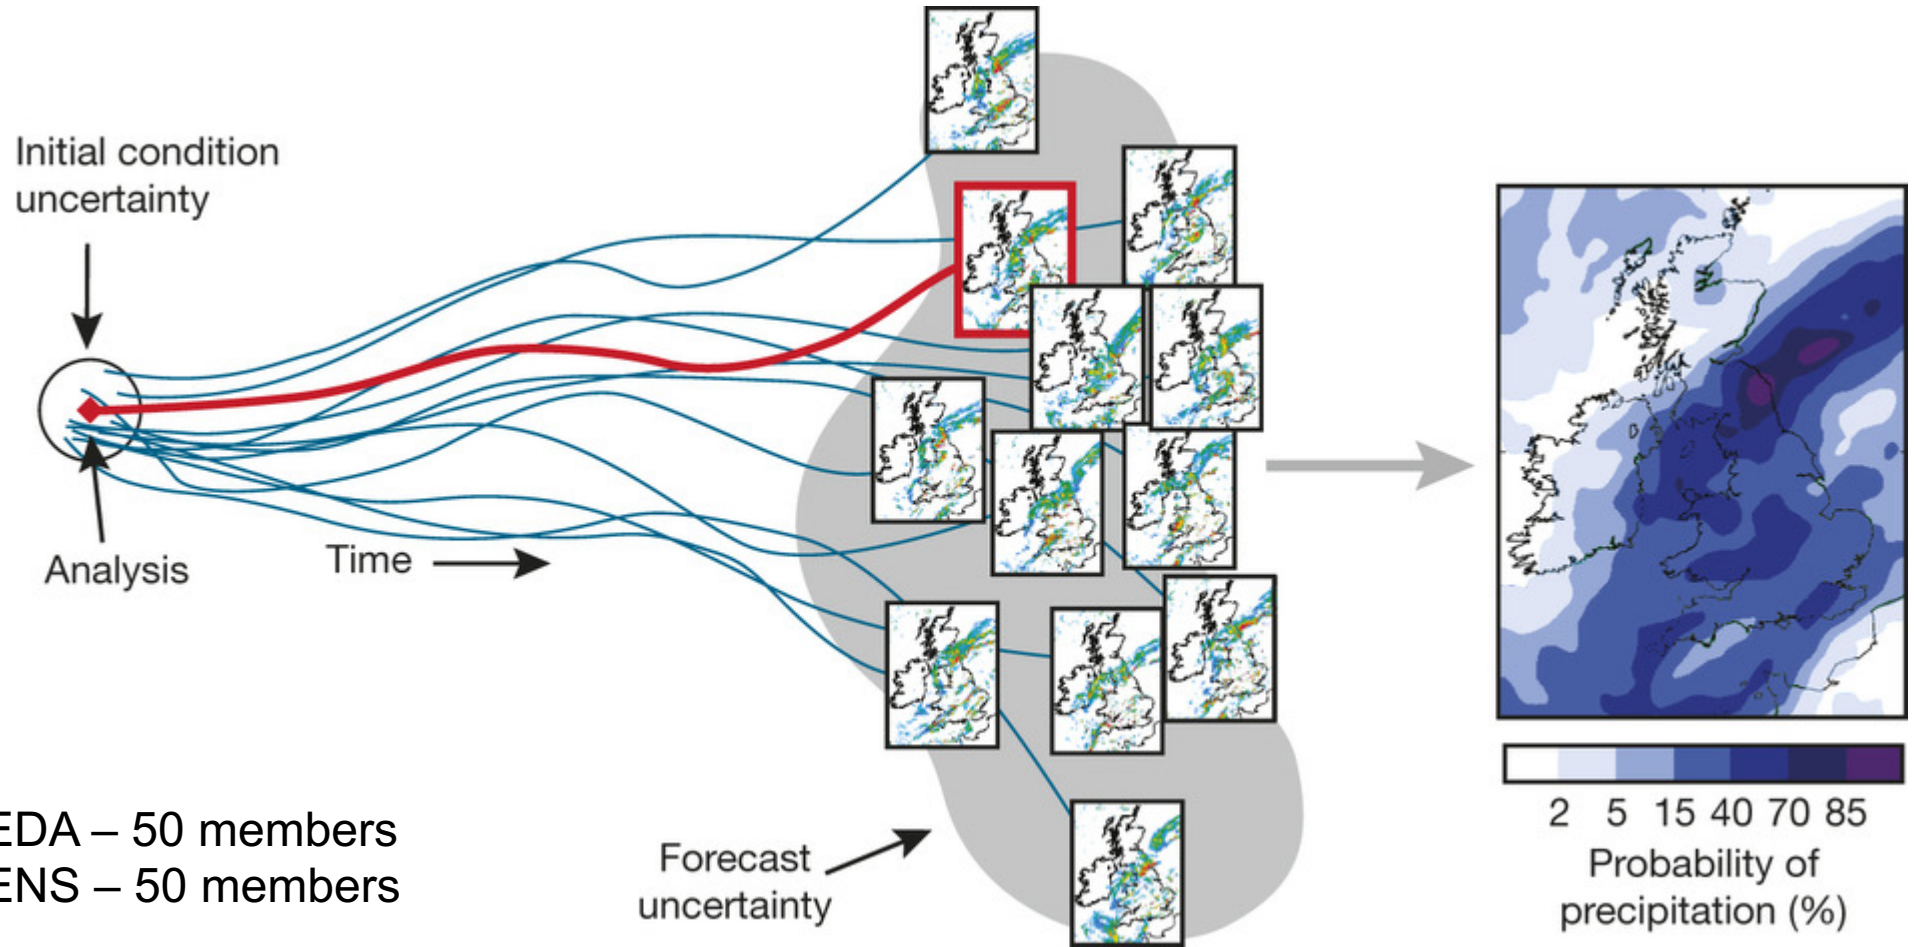
857

The ECMWF Scalability Programme: Progress and Plans

Peter Bauer, Tiago Quintino, Nils Wedi, Antonino Bonanni, Marcin Chrust, Willem Deconinck, Michail Diamantakis, Peter Düben, Stephen English, Johannes Flemming, Paddy Gillies, Ioan Hadade, James Hawkes, Mike Hawkins, Olivier Iffrig, Christian Kühnlein, Michael Lange, Peter Lean, Olivier Marsden, Andreas Müller, Sami Saarinen, Domokos Sarmany, Michael Sleigh, Simon Smart, Piotr Smolarkiewicz, Daniel Thiemert, Giovanni Tumolo, Christian Weihrauch, Cristiano Zanna

February 2020

Ensemble of assimilations and forecasts

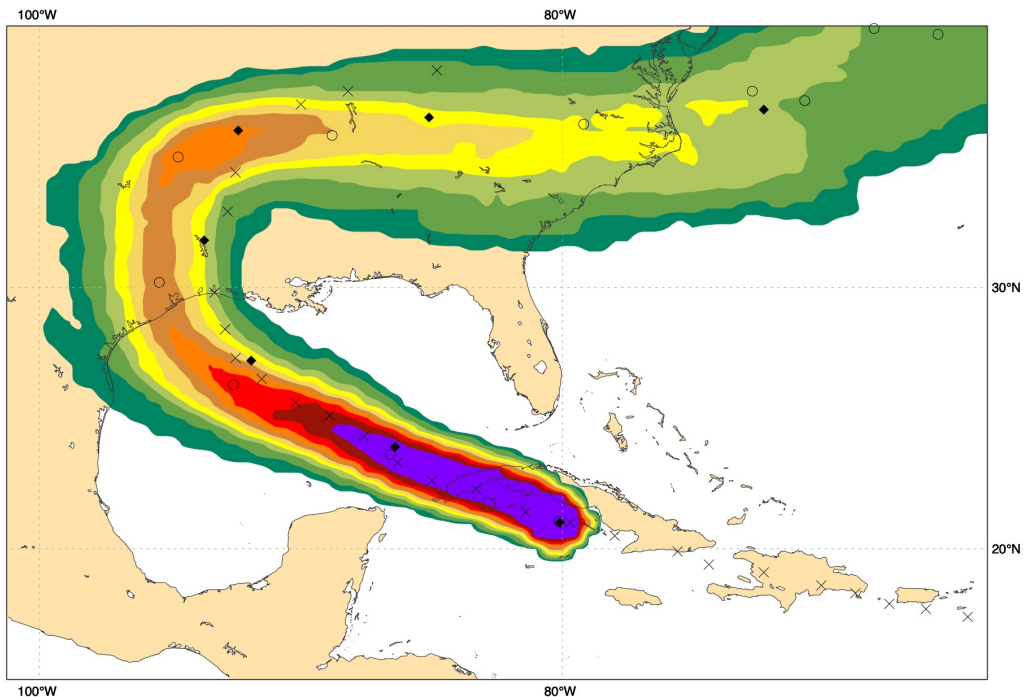


- EDA – 50 members
- ENS – 50 members

Oper, Tco639L91

Date 20200824 12 UTC @ECMWF

Probability that **LAURA** will pass within 120 km radius during the next **240** hours
 tracks: **solid**=HRES; **dot**=Ens Mean [reported minimum central pressure (hPa) **1002**]

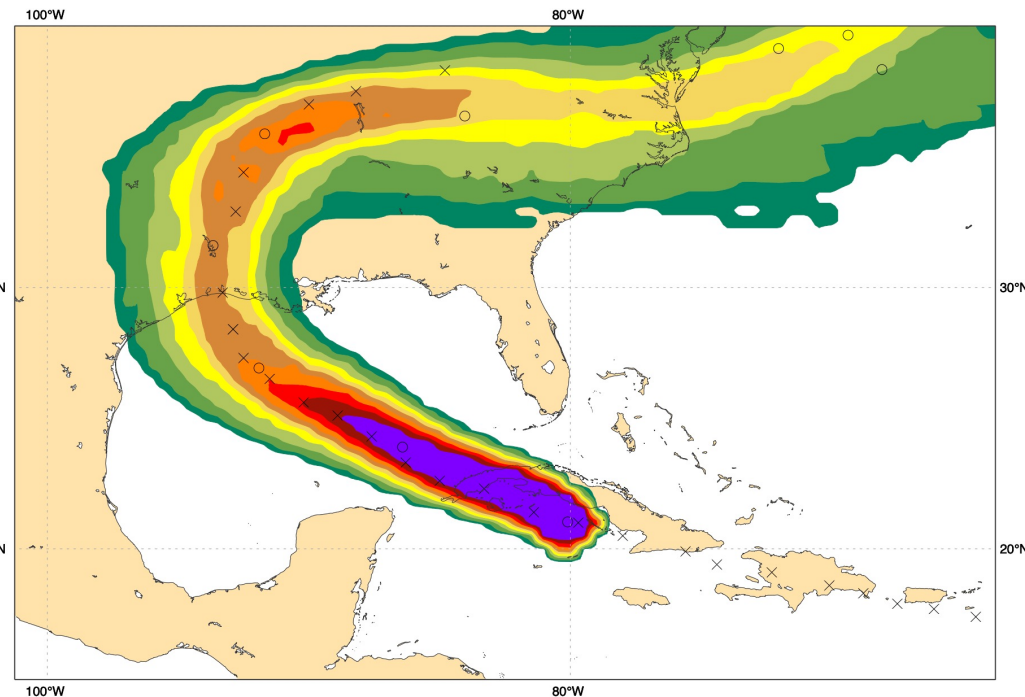


Crosses : observed position
 Circle : ensemble mean
 Diamonds : Oper HRES



Oper, Tco1279L137

Date 20200824 12 UTC @ECMWF


Probability that **LAURA** will pass within 120 km radius during the next **240** hours
 tracks: **solid**=HRES; **dot**=Ens Mean [reported minimum central pressure (hPa) **1002**]




Red dot : observed core pressure
 Solid black line : Oper HRES
 Box plot : ensemble distribution



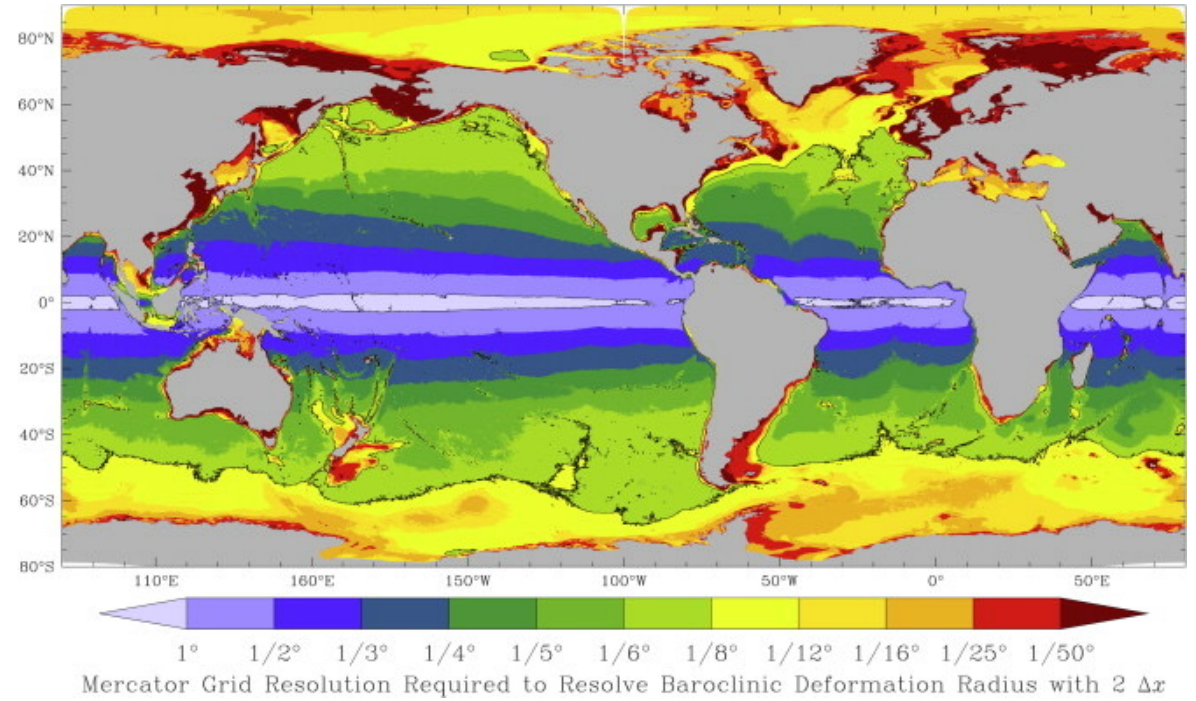
Ocean – Land – Atmosphere – Sea ice



Ocean model - resolution



Hewitt et al. (2017)



Hallberg (2013)

Increased realism in water cycle reservoir representation at 1km

Ioan Hadade, Gabriele Arduini, Souhail Boussetta, Joey McNorton, Margarita Choulga, Gianpaolo Balsamo, et al.

The **Offline Surface** Modelling (OSM) **increased performance** allows to run the surface at **1km at ECMWF**

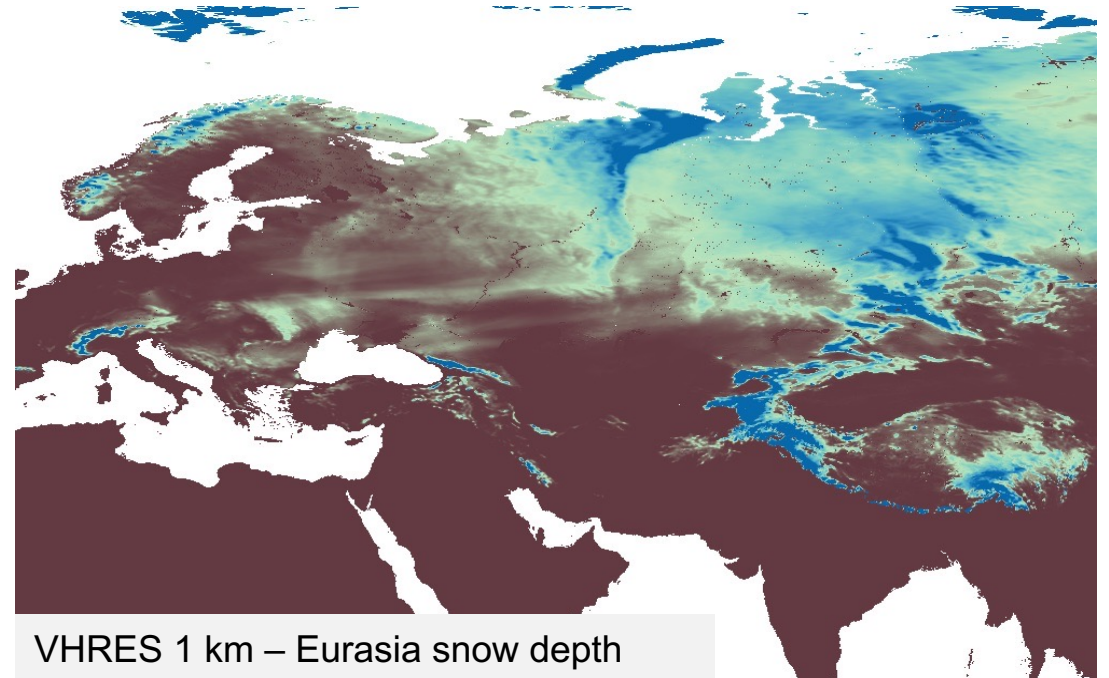
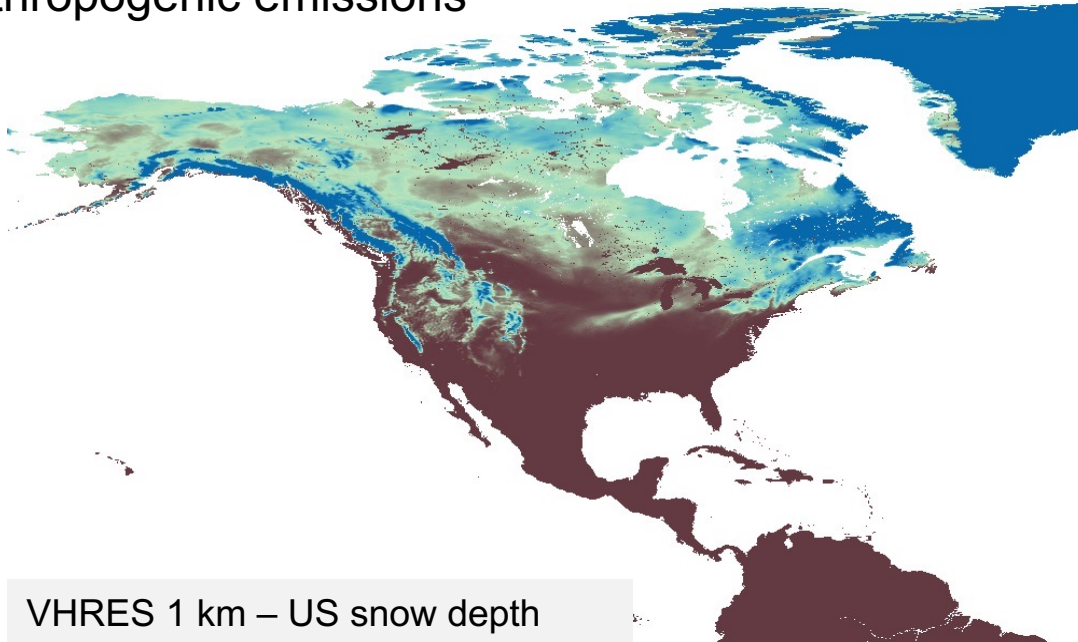
Increased realism of

- Land use and land cover (use of ESA-CCI)
- Coastal areas and lakes (use of GSWE)
- Snow over orography & catchment hydrology
- Urban areas

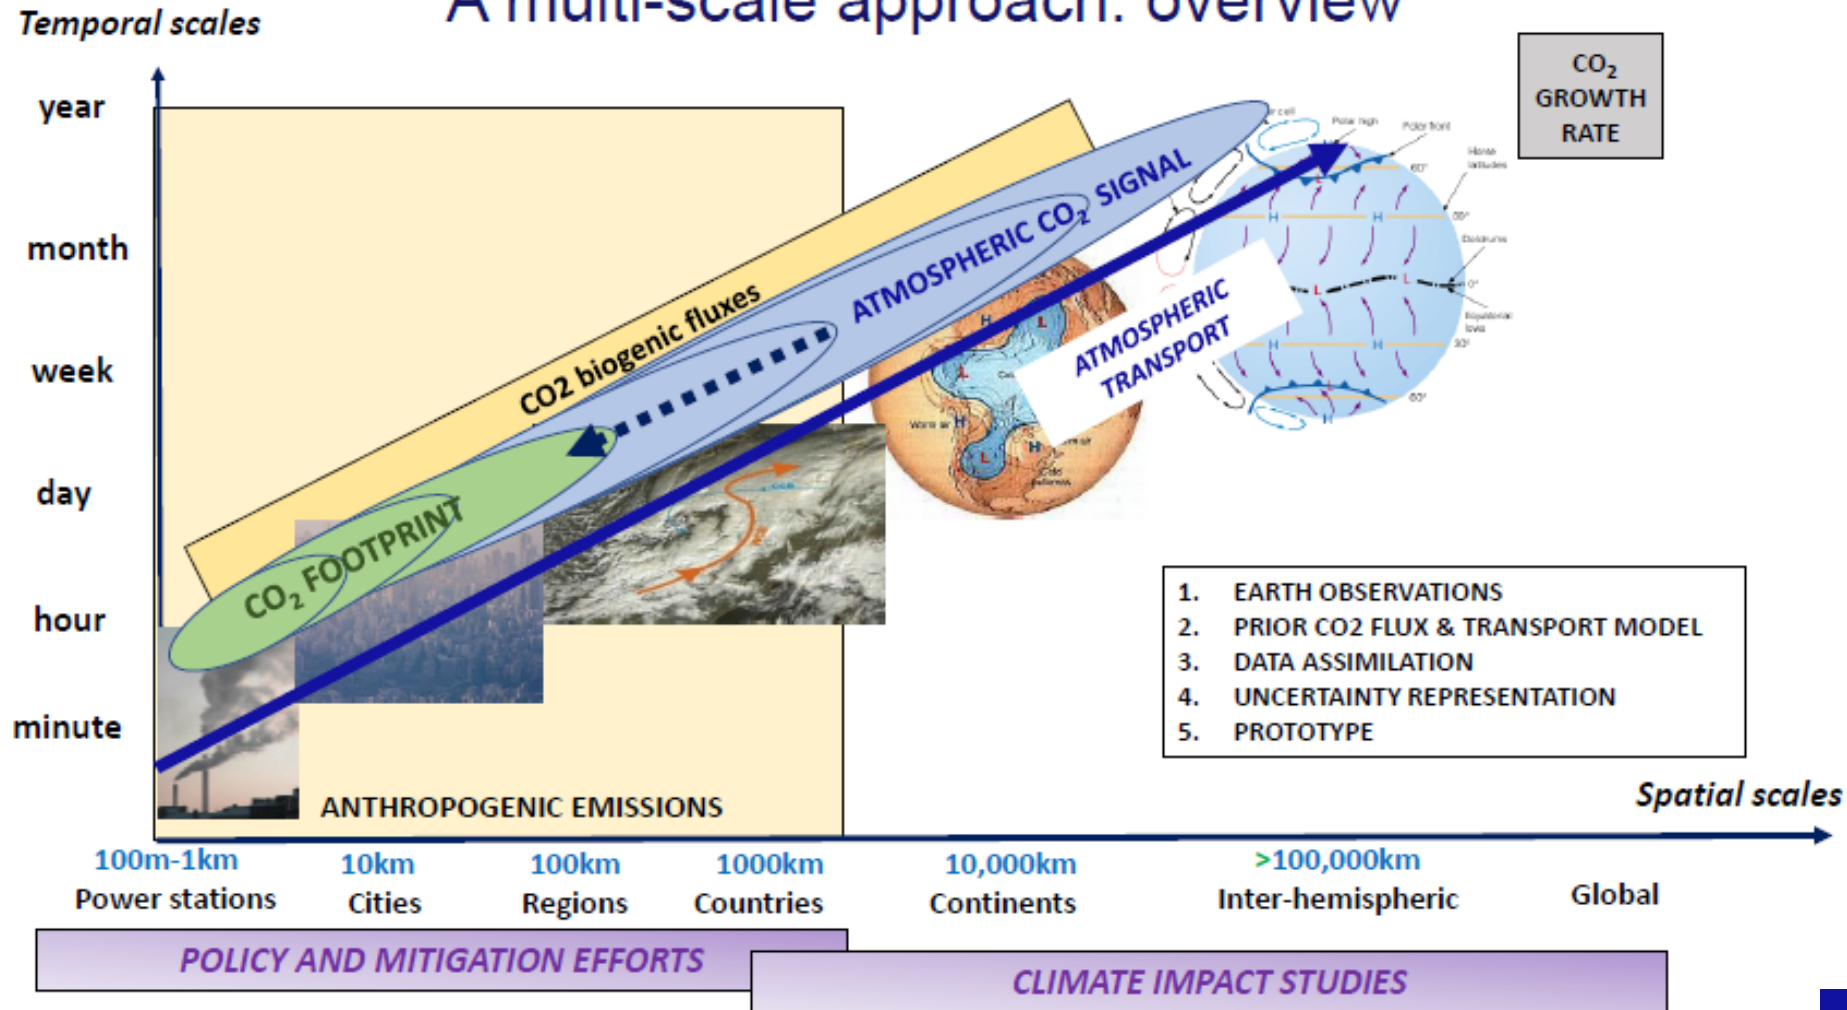
Prerequisite for improved analysis of

- skin temperature
- anthropogenic emissions

Resolution	Configuration	Performance SYPD
9km (HRES & ERA5Land)	TCo1279	~ 8
1km (VHRES)	TCo7999	~ 1



A multi-scale approach: overview



Key aspects for Numerical methods:


- Horizontal and vertical resolution at the scale of emission sources (100m - 1km)
- Transport uncertainty and tracer conservation
- Cost of many advected tracers and shape preservation (e.g. monotonicity, positivity, thresholds)



Improvements to NWP systems in the Stratosphere: better reanalysis & monitoring

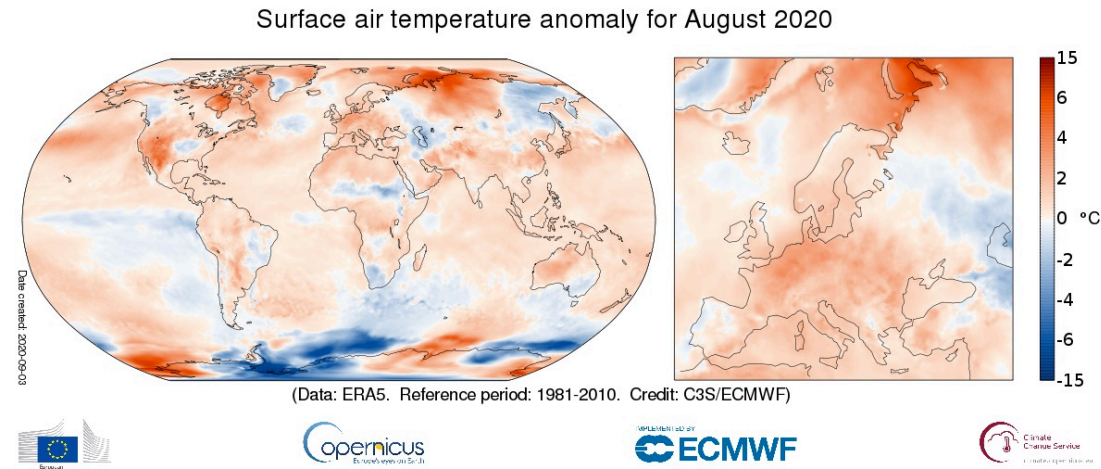
Much better representation of Sudden Stratospheric Warming events, due to changes in the Semi-Lagrangian scheme
(*Diamantakis, 2014*)

NH winter SSWs



Standard deviation of MW radiances observed vs simulated temperature fields of ERA-Interim (blue) and ERA5 (red) using satellite channel (noaa15) peaking around 5hpa.

T. McNally, A. Simmons



Towards increasing realism with storm-scale resolving simulations

- The ECMWF Cubic Octahedral (**ECO1280**) Nature Run (*Ross Hoffman, Tanya Peevey, S. Malardel, ...*)
 - <https://www.cira.colostate.edu/imagery-data/ecmwf-nature-run/>
- **DYAMOND**: the DYnamics of the Atmospheric general circulation Modelled On Non-hydrostatic Domains
 - Stevens et al, 2019 <https://progearthplanetsci.springeropen.com/articles/10.1186/s40645-019-0304-z>
- **DYAMOND II** (ongoing)
 - Simulation period Jan – Mar 2020 (40 days), shadowing EURECA campaign, including coupled simulations
 - IFS at 4km coupled to NEMO
 - IFS at 4km coupled to FESOM2
 - Additional: NH-IFS with Arome and IFS physics at 2.5km
- Global simulations of the atmosphere at 1.4 km grid-spacing with the Integrated Forecasting System (*Dueben, Wedi, Saarinen, Zeman, JMSJ 2020*)
 - First runs on PizDaint (Europe's biggest computer), NH-IFS and H-IFS comparisons

INCITE awards computing time on the 2nd largest supercomputer in the world (Summit, top500 June 2020) at the Oak Ridge Leadership Computing Facility (OLCF), Oakridge, Tennessee, US




~4600 nodes:
2 x ibm_power9 (42 cores)
6 x NVIDIA V100 GPUs
512 GB DDR4 memory
1.6 TB NVMe
96 GB HBM2 (GPU only)

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

Motivations:

- Push accelerator readiness
- Status and capability of the IFS model
- Comparison of dynamical cores
- Explicitly simulating deep convection
- Explore impact on longer time scales
- Support for OSSEs

3-hourly accumulated radiative fluxes at the top of the atmosphere



1.4 km

9 km

9 km w/o deep


A baseline for global weather and climate simulations at 1 km resolution

Nils P. Wedi¹, Inna Polichtchouk, Peter Dueben¹, Valentine G. Anantharaj², Peter Bauer¹, Souhail Boussetta¹, Philip Browne¹, Willem Deconinck¹, Wayne Gaudin³, Ioan Hadade¹, Sam Hatfield¹, Olivier Iffrig¹, Philippe Lopez¹, Pedro Maciel¹, Andreas Mueller¹, Sami Saarinen¹, Irina Sandu¹, Tiago Quintino¹, Frederic Vitart¹


(JAMES, 2020, avail online)

1.4 km

2018110100+48h




Meteosat-8, Meteosat-11 and GOES-15
2018110300




Philippe Lopez & Cristina Lupu

9 km

2018110100+48h




Meteosat-8, Meteosat-11 and GOES-15
2018110300




Philippe Lopez & Cristina Lupu

9 km w/o deep

2018110100+48h

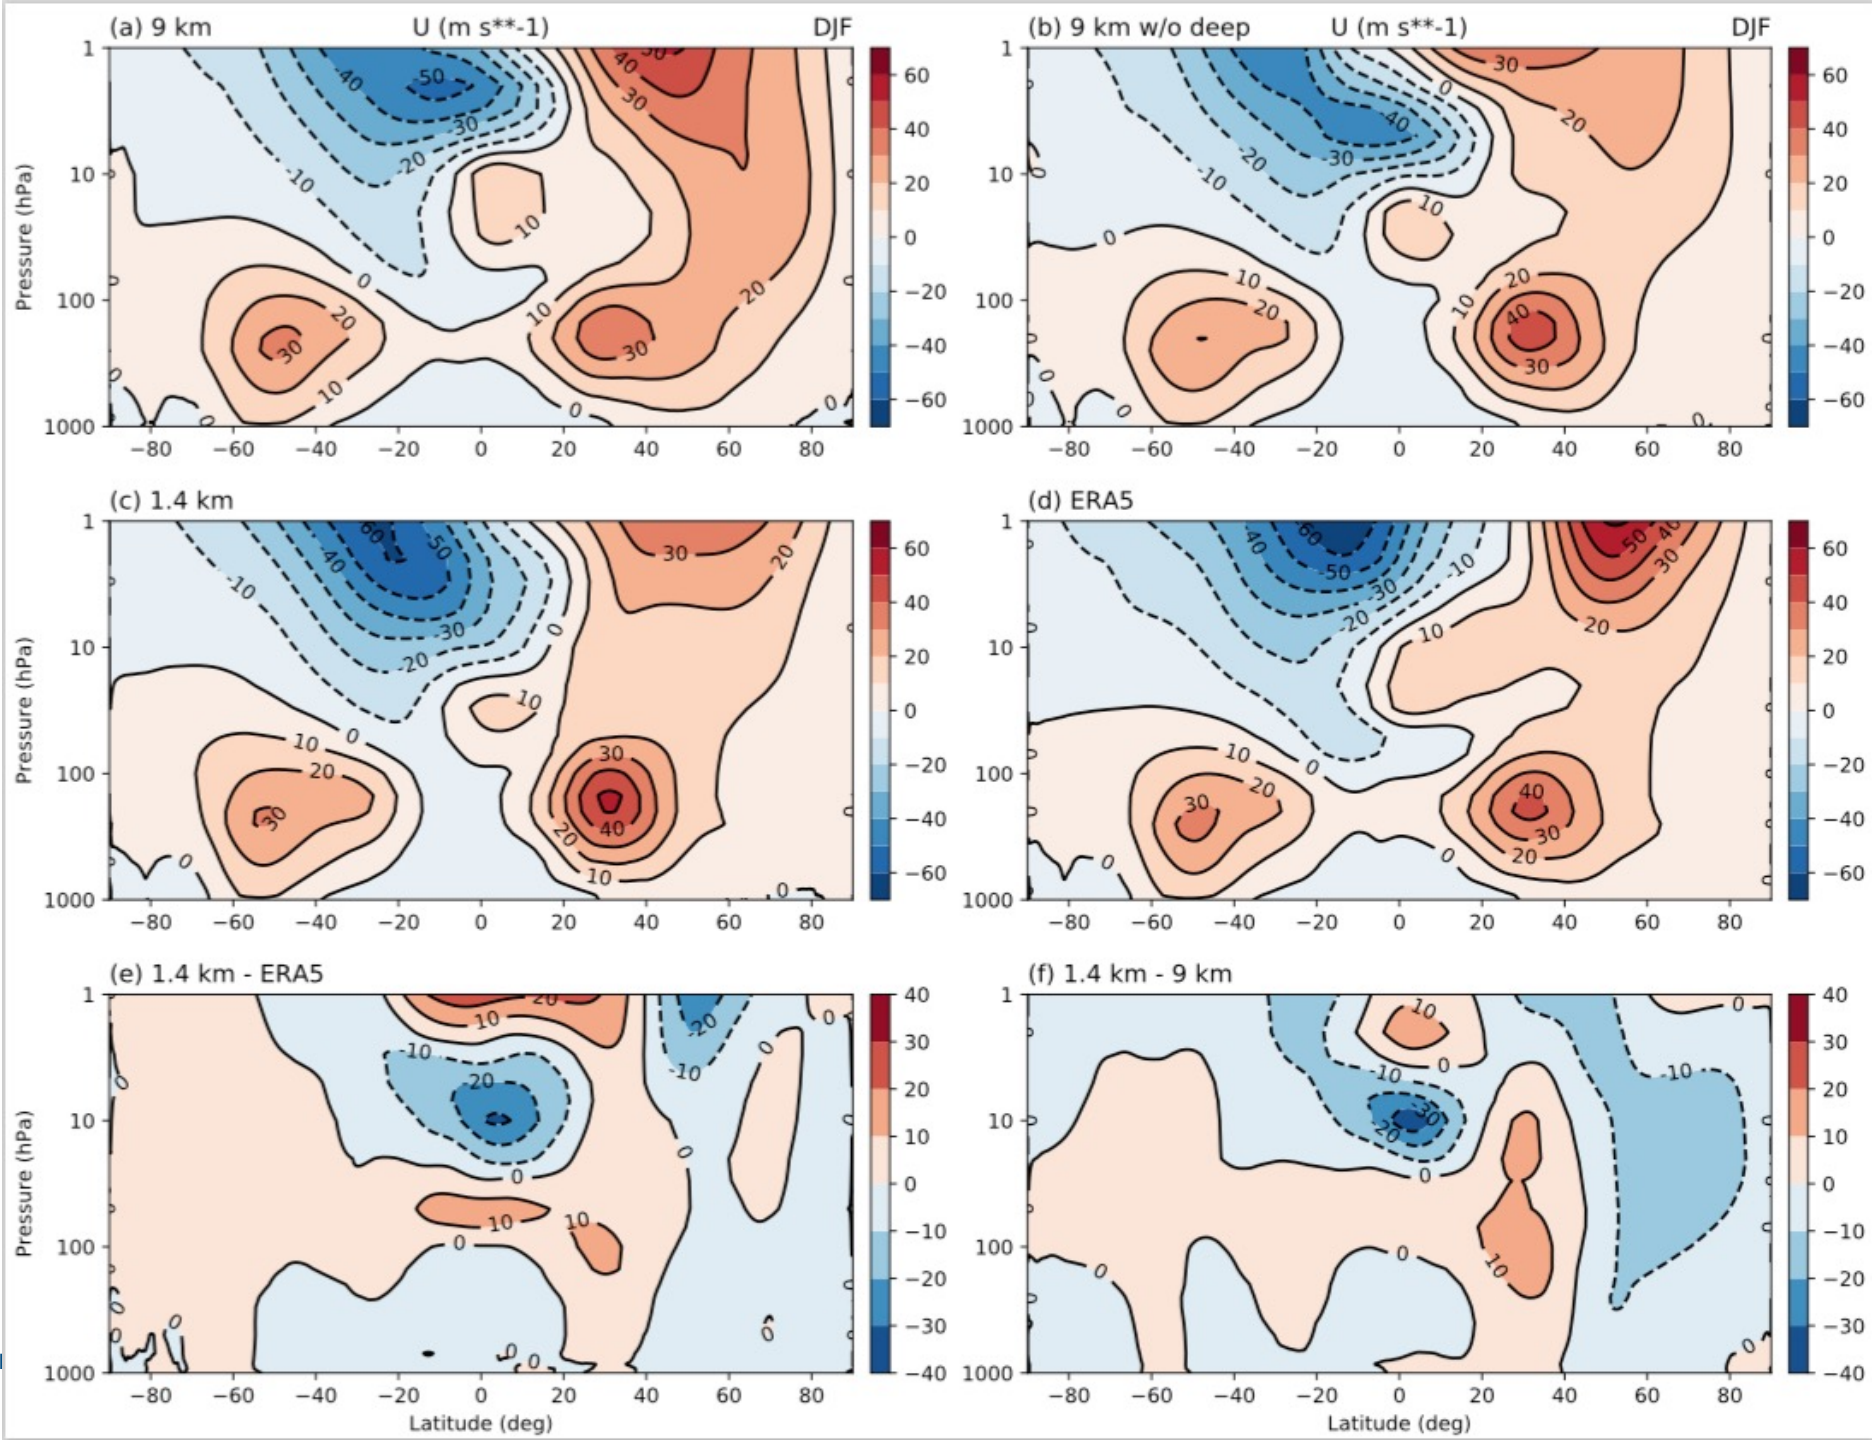


Meteosat-8, Meteosat-11 and GOES-15
2018110300

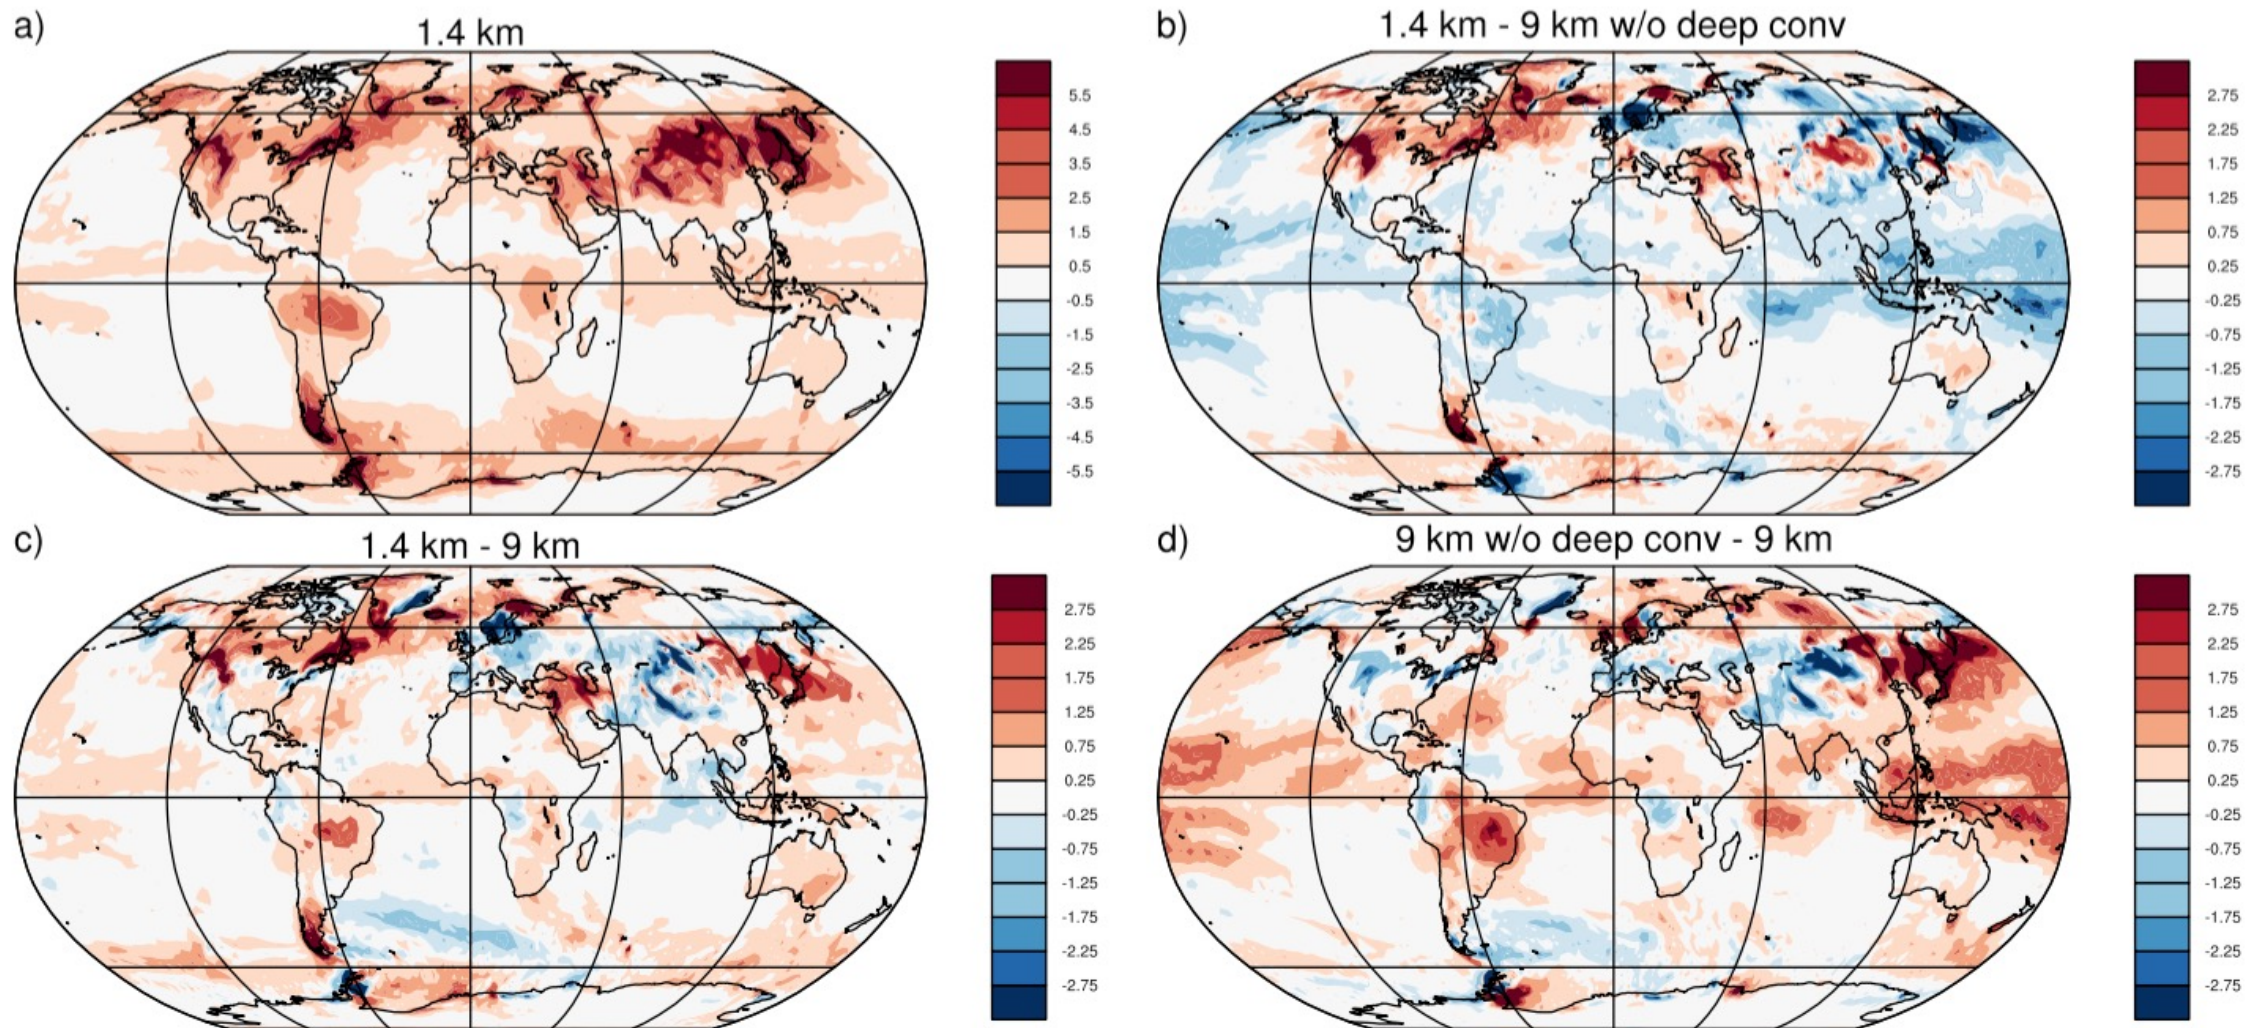


Philippe Lopez & Cristina Lupu

Zonal-mean zonal wind



And what you don't see in the satellite pics ...



Zonal-mean absolute gravity wave momentum flux [mPa], computed from the total wave numbers 42-1279 for November 2018 at 50 hPa enhanced in convectively active regions, visibly much stronger with explicitly simulated convection at 9 km grid spacing.

Summary

- The spectral transform method continues to be competitive for global models at km-scales
- A first digital twin prototype has been presented with a global seasonal simulation at 1.4 km grid-spacing on Summit
- Advancing numerical methods continues to be a key contribution towards time and energy efficiency and towards realising the ambitious goal of routine global km-scale data assimilation and prediction of the coupled Earth System
- Big data handling, unsupervised learning, and near-real time tailored impact sector interaction (e.g. Energy,Health,Hydrology,Biodiversity) forms an integral part of future km-scale coupled model development

Machine learning at ECMWF

P. Dueben




Figure copied from www.wikipedia.org

- Neural Networks can learn from input/output pairs to emulate a non-linear process.
- Neurons have weighted connections to each other and the weights are trained to produce the optimal results.
- There are plenty of model components in Earth System models that show non-linear behaviour that can serve as applications for neural networks.

Observation processing:

- Real-time quality control of observations
- Detection of unrealistic weather situations and discrepancies between products
- Bias correction
- Feature detection to reduce data volume
- ...

Numerical modelling:

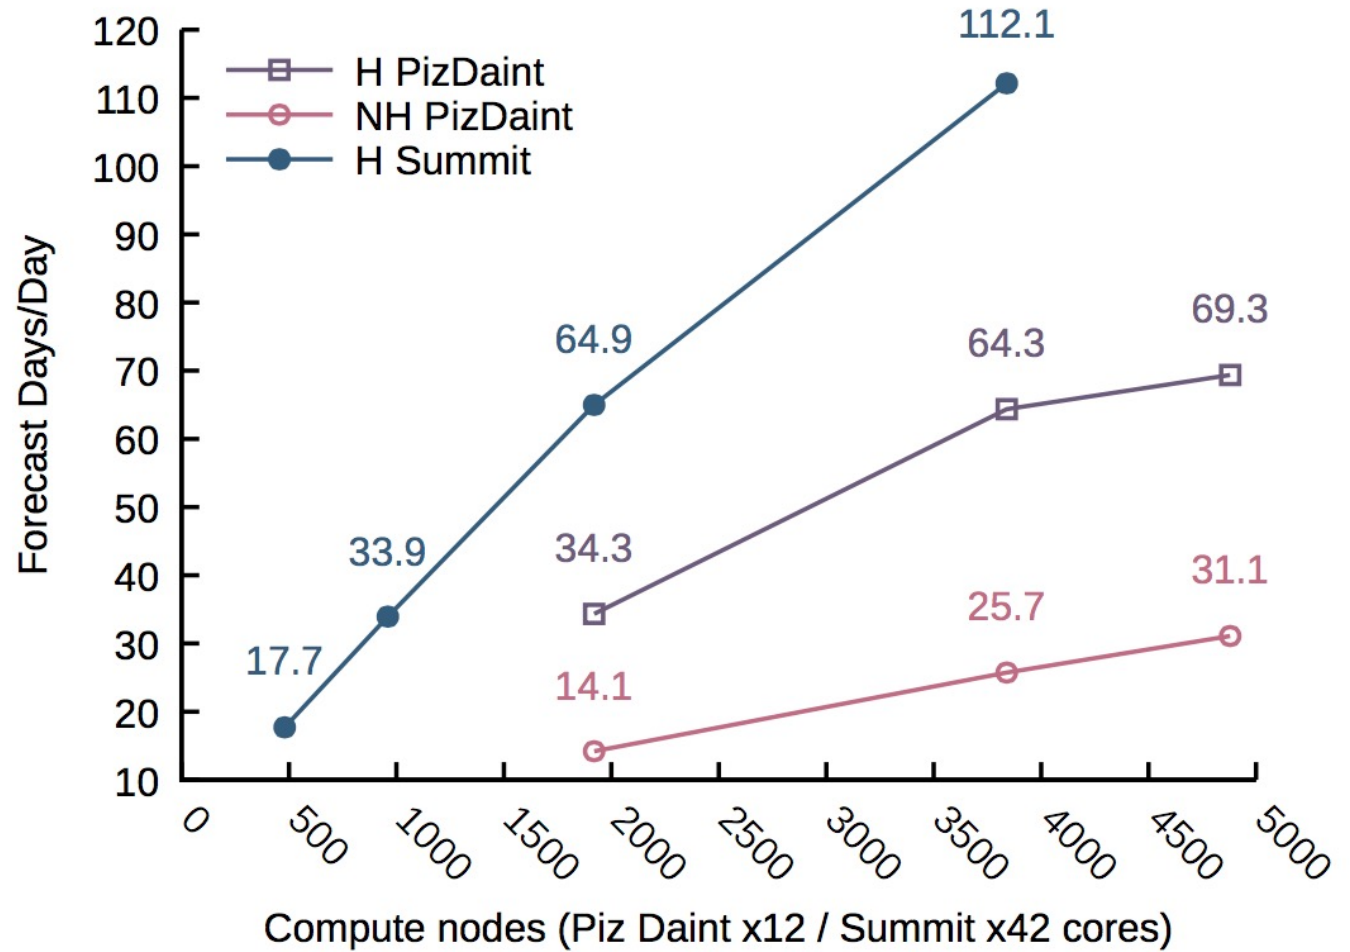
- Emulation for efficiency
- Emulation for portability
- Emulation for generation of TL/AD code
- Estimation of model bias in data assimilation
- Improvement of parametrisation schemes
- ...

Post-processing:

- Real time adjustments
- Bias correction
- Local Downscaling
- Feature detection
- Uncertainty quantification
- Error correction for seasonal predictions
- ...



Strong scaling at 1.4 km



This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

GPU Spectral transform dwarf

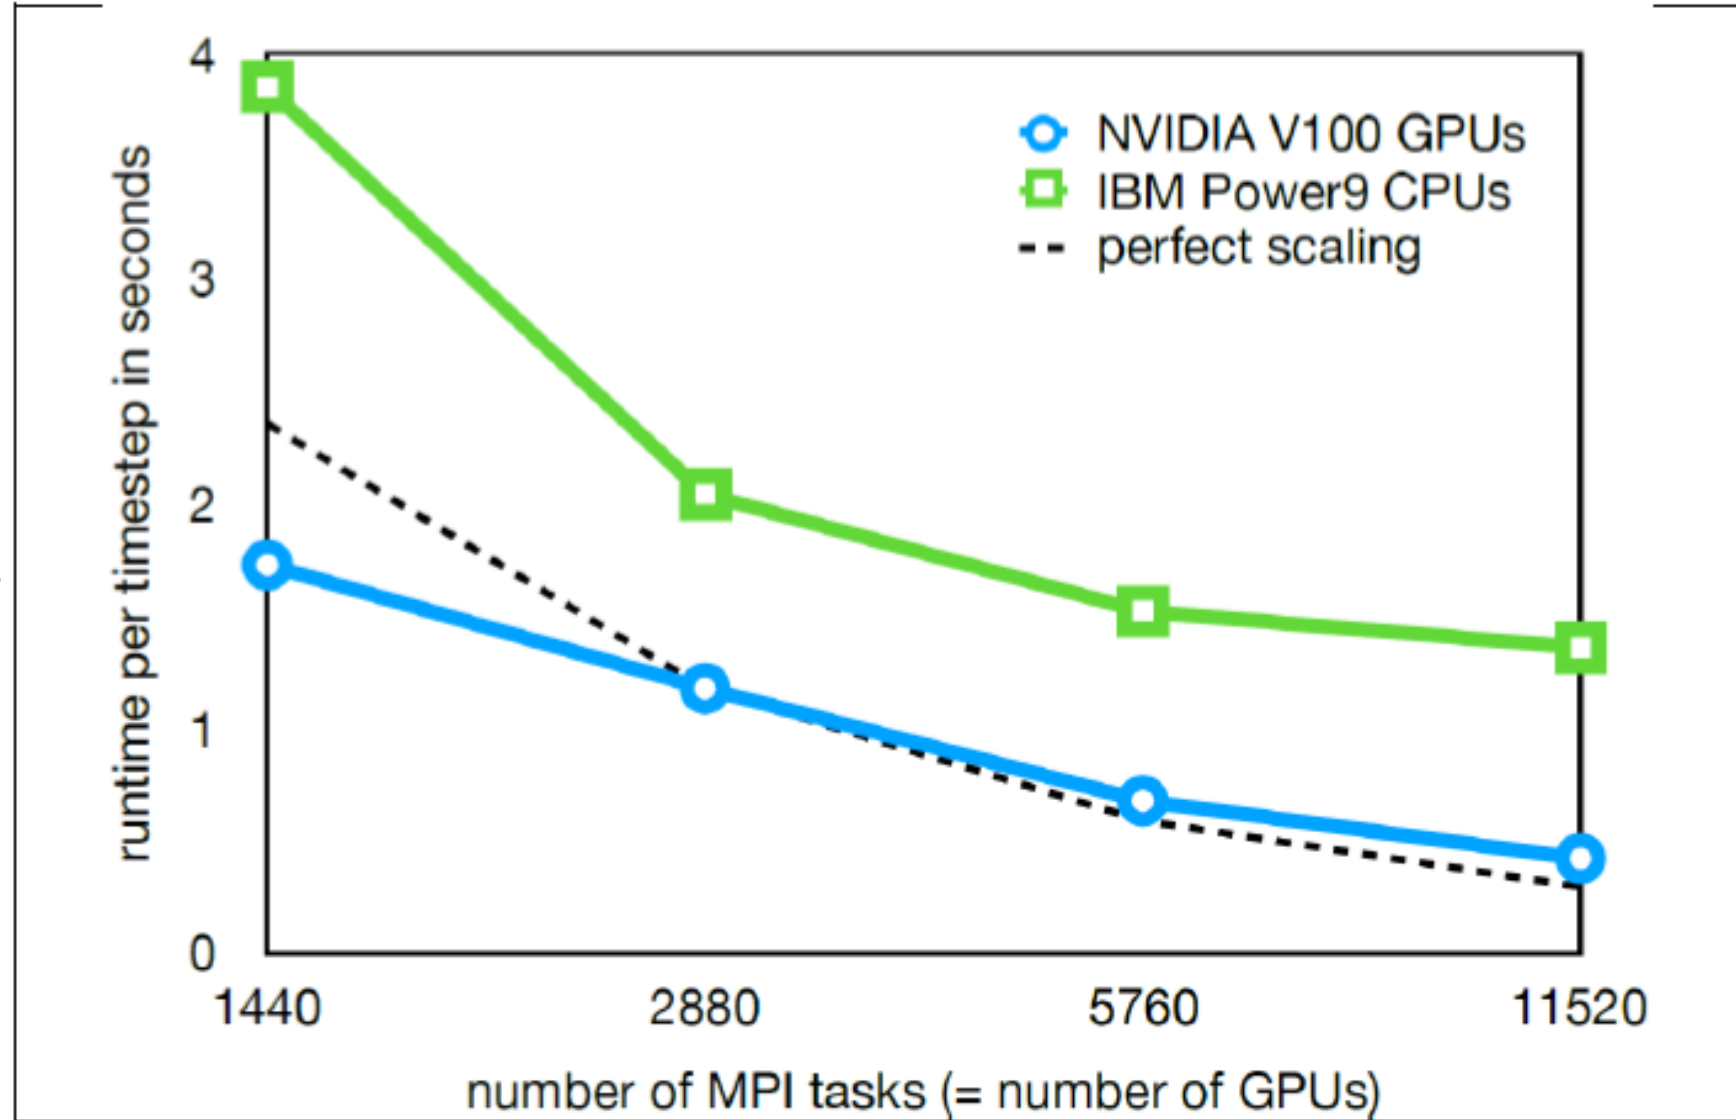
Andreas Mueller, based on initial work by Alan Gray and recently also Wayne Gaudin (NVIDIA)



ESCAPE 2

Choices: PGI, GNU, XL
MPI/OpenMP/OpenACC/cuLibraries

At 2.9km resolution, less than 1s per time-step fits operational needs.



This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.