Towards a Digital Twin of the Earth System

Nils P. Wedi
European Centre for Medium-Range Weather Forecasts (ECMWF)

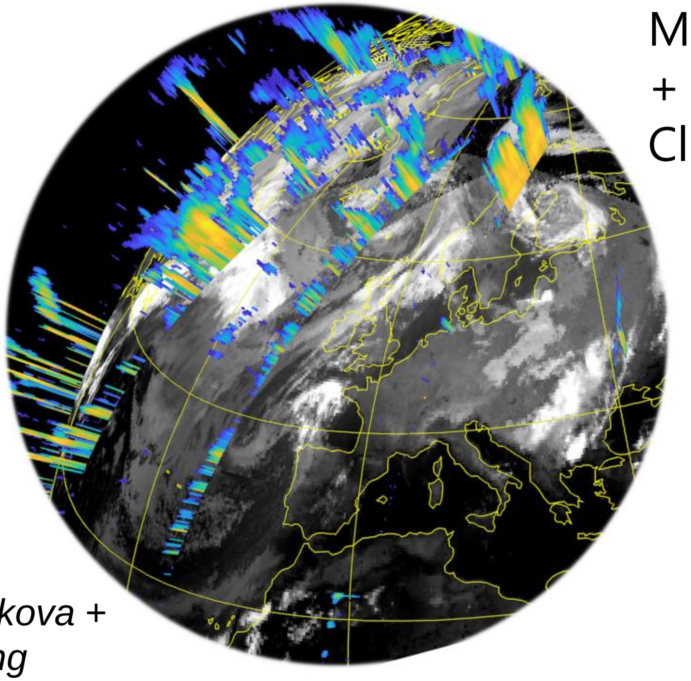

MODIS aqua infrared channel
+ cloud radar cross sections from
CloudSat

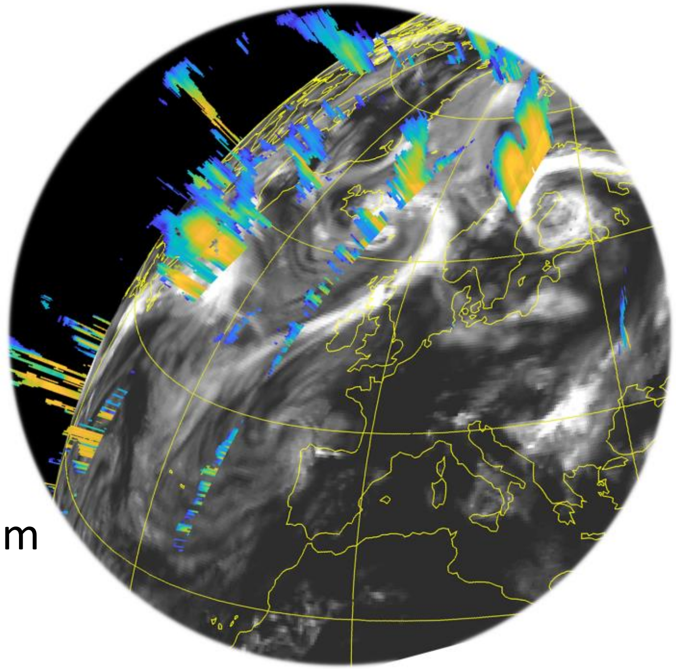

Simulated satellite image from
the IFS model with analysed
cross-section cloud profiles

Marta Janiskova +
Mark Fielding

Outline

- A brief review of the evolving model component(s) of the Integrated Forecasting System (IFS)
- Increasing complexity of the evolving Earth-System to improve data assimilation and predictions at different spatio-temporal scales (diurnal to seasonal)
- Performance and Portability
- 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

Schematic description of the *spectral transform method* in the ECMWF IFS model

Eliassen (1970); Orszag, (1970), … Machenhauer, Bourke, Robert, Ritchie, Williamson, Swarztrauber, Jarraud, Simmons, Burridge, Hoskins, Courtier, Geleyn, Temperton, Hortal, Hamrud, …

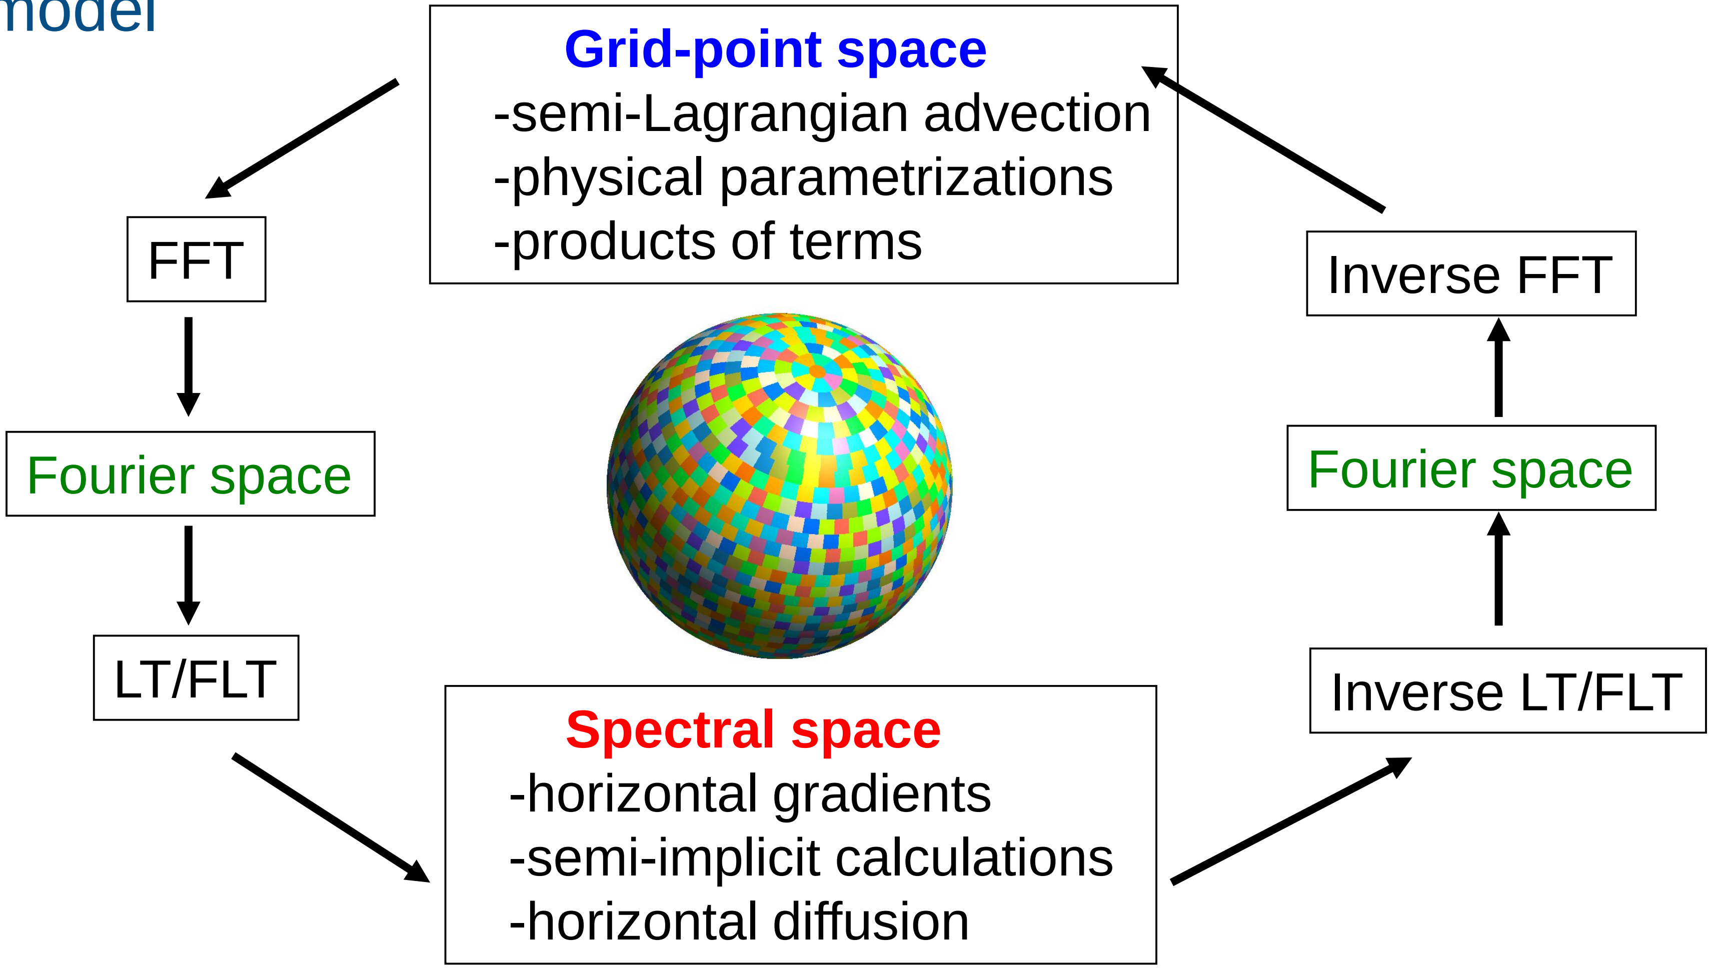

Grid-point space
- -semi-Lagrangian advection
- -physical parametrizations
- -products of terms

FFT

Fourier space

LT/FLT

Spectral space
- -horizontal gradients
- -semi-implicit calculations
- -horizontal diffusion

Inverse LT/FLT

Fourier space

Inverse FFT

FFT: Fast Fourier Transform, LT/FLT: Legendre Transform

(Wedi et al, 2013, 2014, 2015)

Global spectral methods

Pros	Cons
• Global connectivity and direct access to global scales (and spectra) • Multigrid options and direct aliasing control (truncation) *(see also talk by M. Ujiie)* • Effective compression of data • Efficient direct semi-implicit solvers (potentially suitable for time-parallel schemes, *see talk by M. Schreiber*) • Large timestep • Fast Legendre transform and high floating point rates in GPU adaptations • Single precision and potential for half precision hardware use (*see talk by S. Hatfield*) • Easy to implement diffusion or global fixers • Use of realistic topography without excessively filtering smaller scales	• Spectral ripples (resulting in spurious time mean fields) • Best suited for global hydrostatic equations • Non-hydrostatic with steep slopes complicate issues with direct solvers (*but see talks by F. Voitus and P. Termonia*) • MPI – parallel communication overhead • Energy and strong scaling efficiency (due to large data movement every timestep) • Formal conservation (often combined with semi-Lagrangian requires global fixers) • Questionable for fully non-hydrostatic scales where dx ~ dz

Accurate representation of orography

- Differences in orography (and associated filters applied) have a significant impact on surface stresses and circulation patterns on weather & climate models.

Elvidge et al 2019

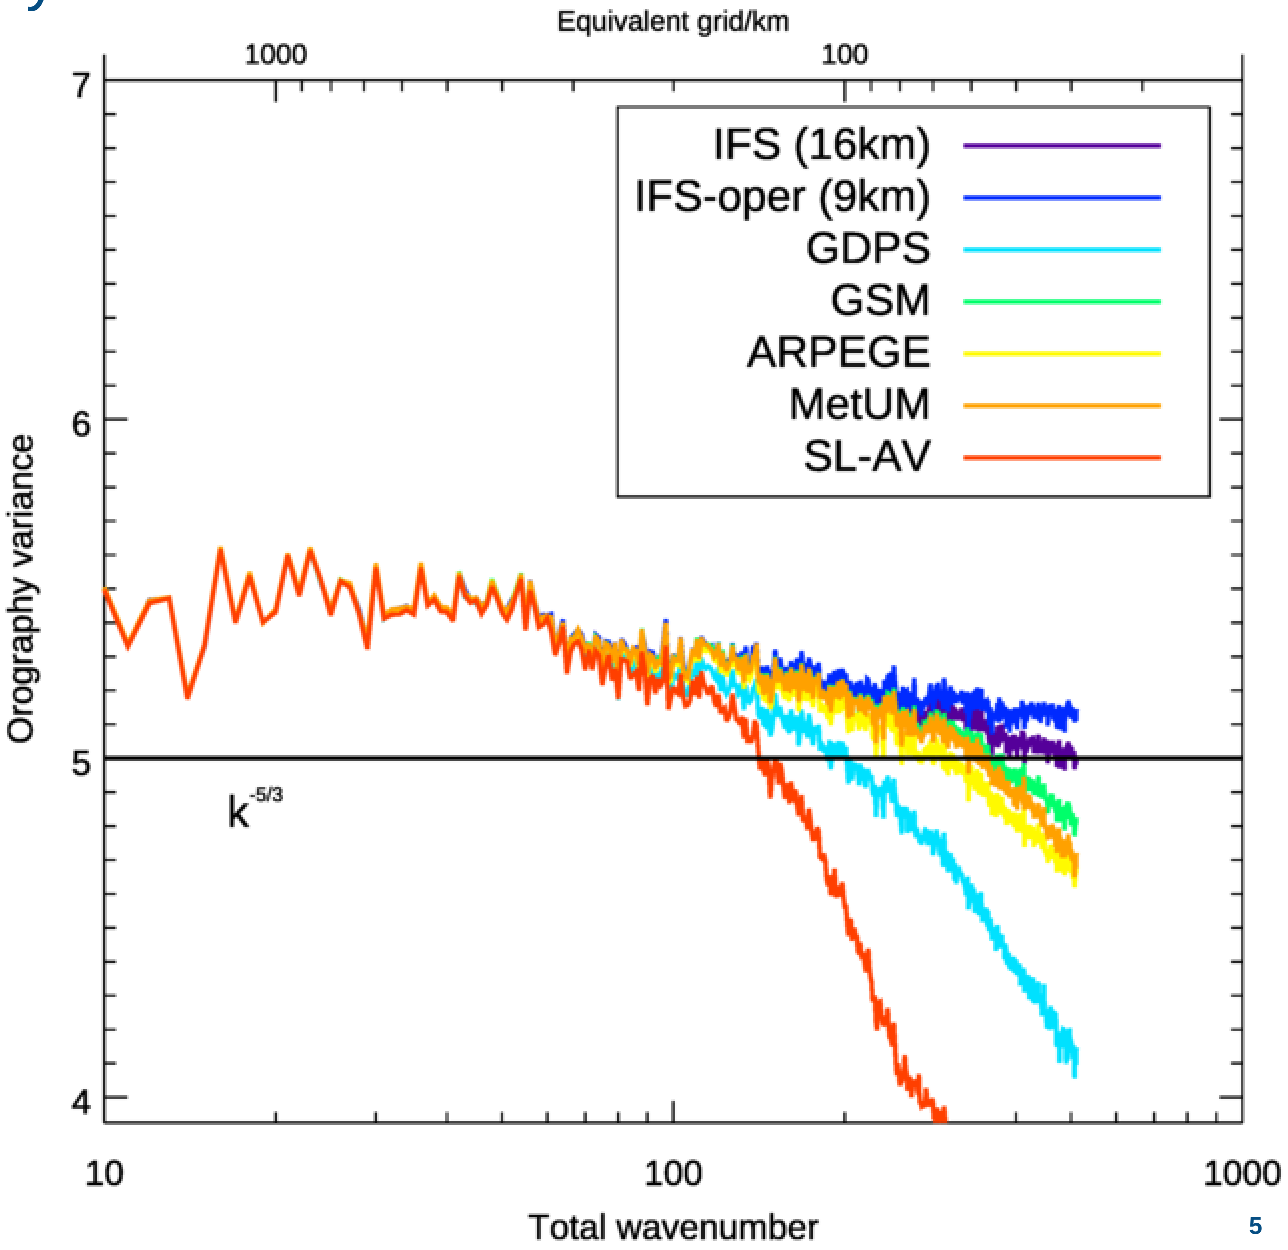

Computational power drives spatial resolution

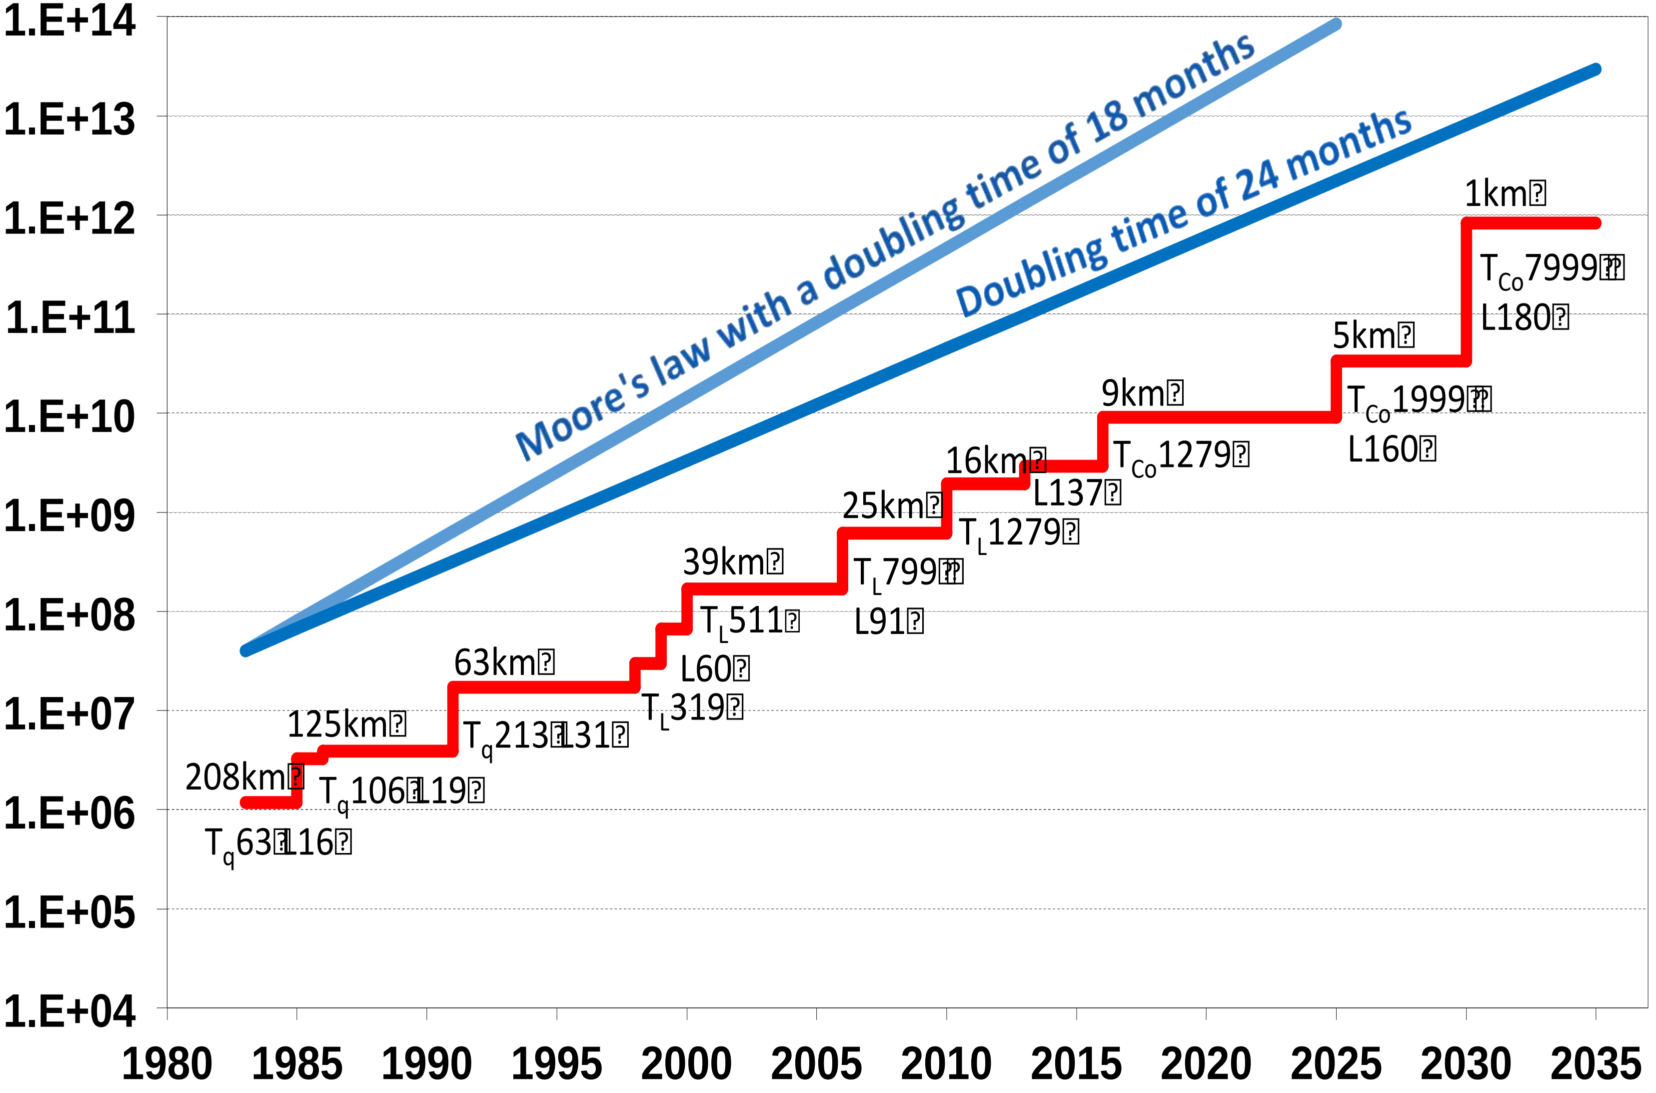

(Schulthess et al, 2019)

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

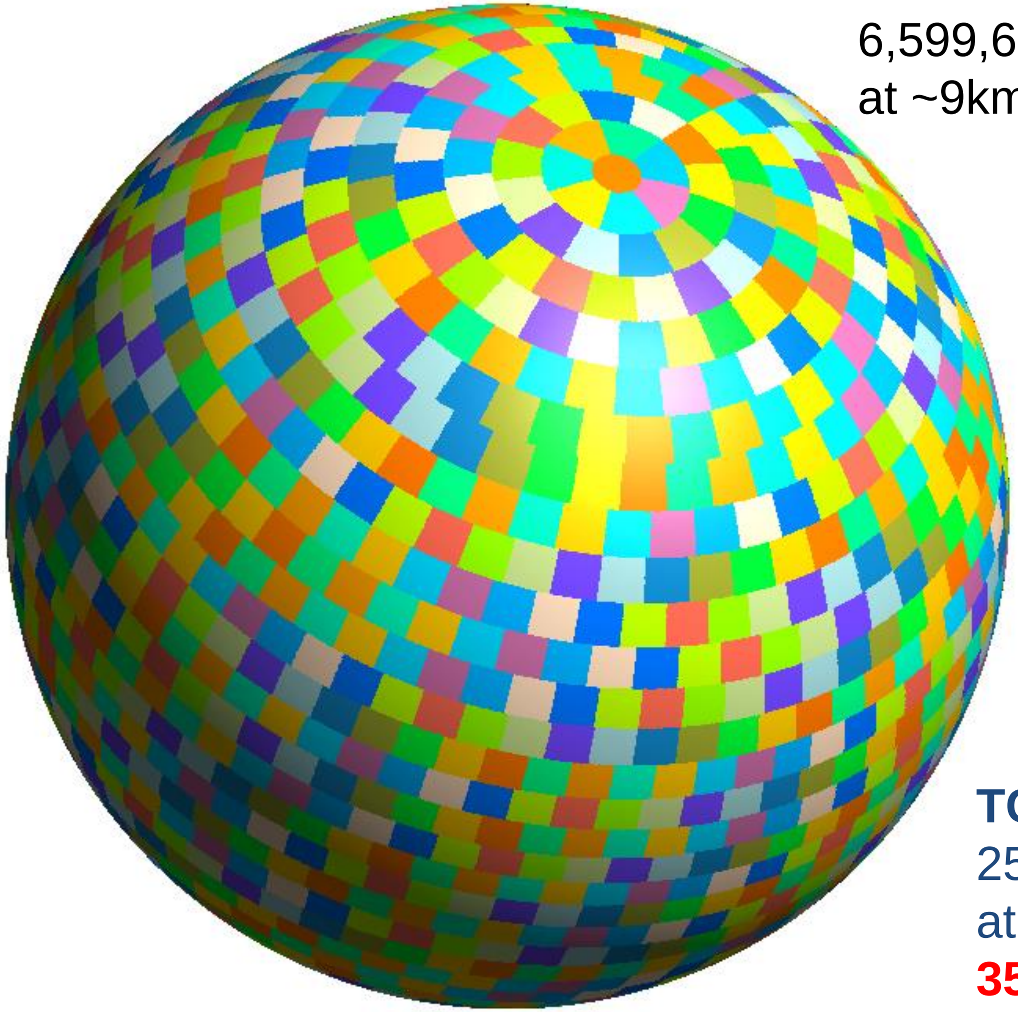

6,599,680 points x 137 levels x 10 vars
at ~9km ~ **9 Billion points or ~100TB/day**

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

A cubic octahedral grid

A further ~20% reduction in gridpoints
=> ~50% less points compared to full grid

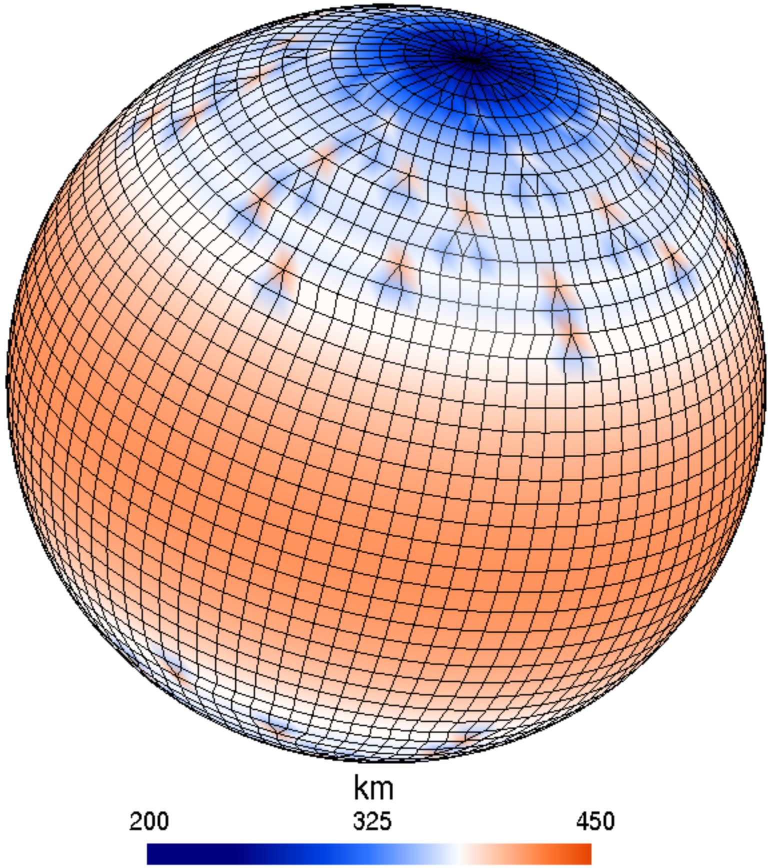

N24 reduced Gaussian grid

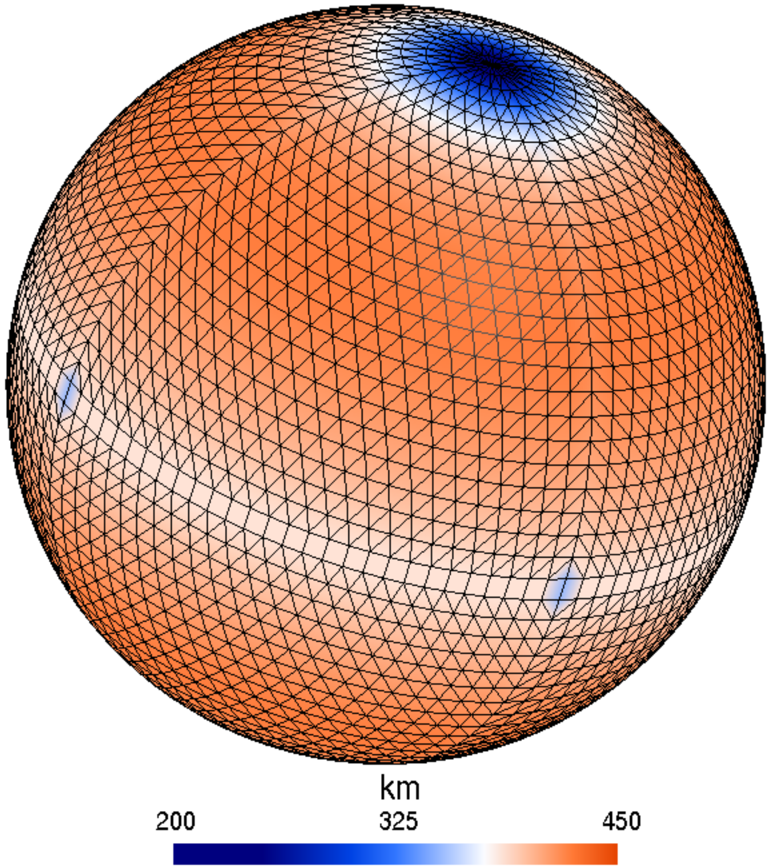

N24 octahedral Gaussian grid

(Wedi et al, 2015)

Atlas: a library for NWP and climate modelling

Funded by the
European Union

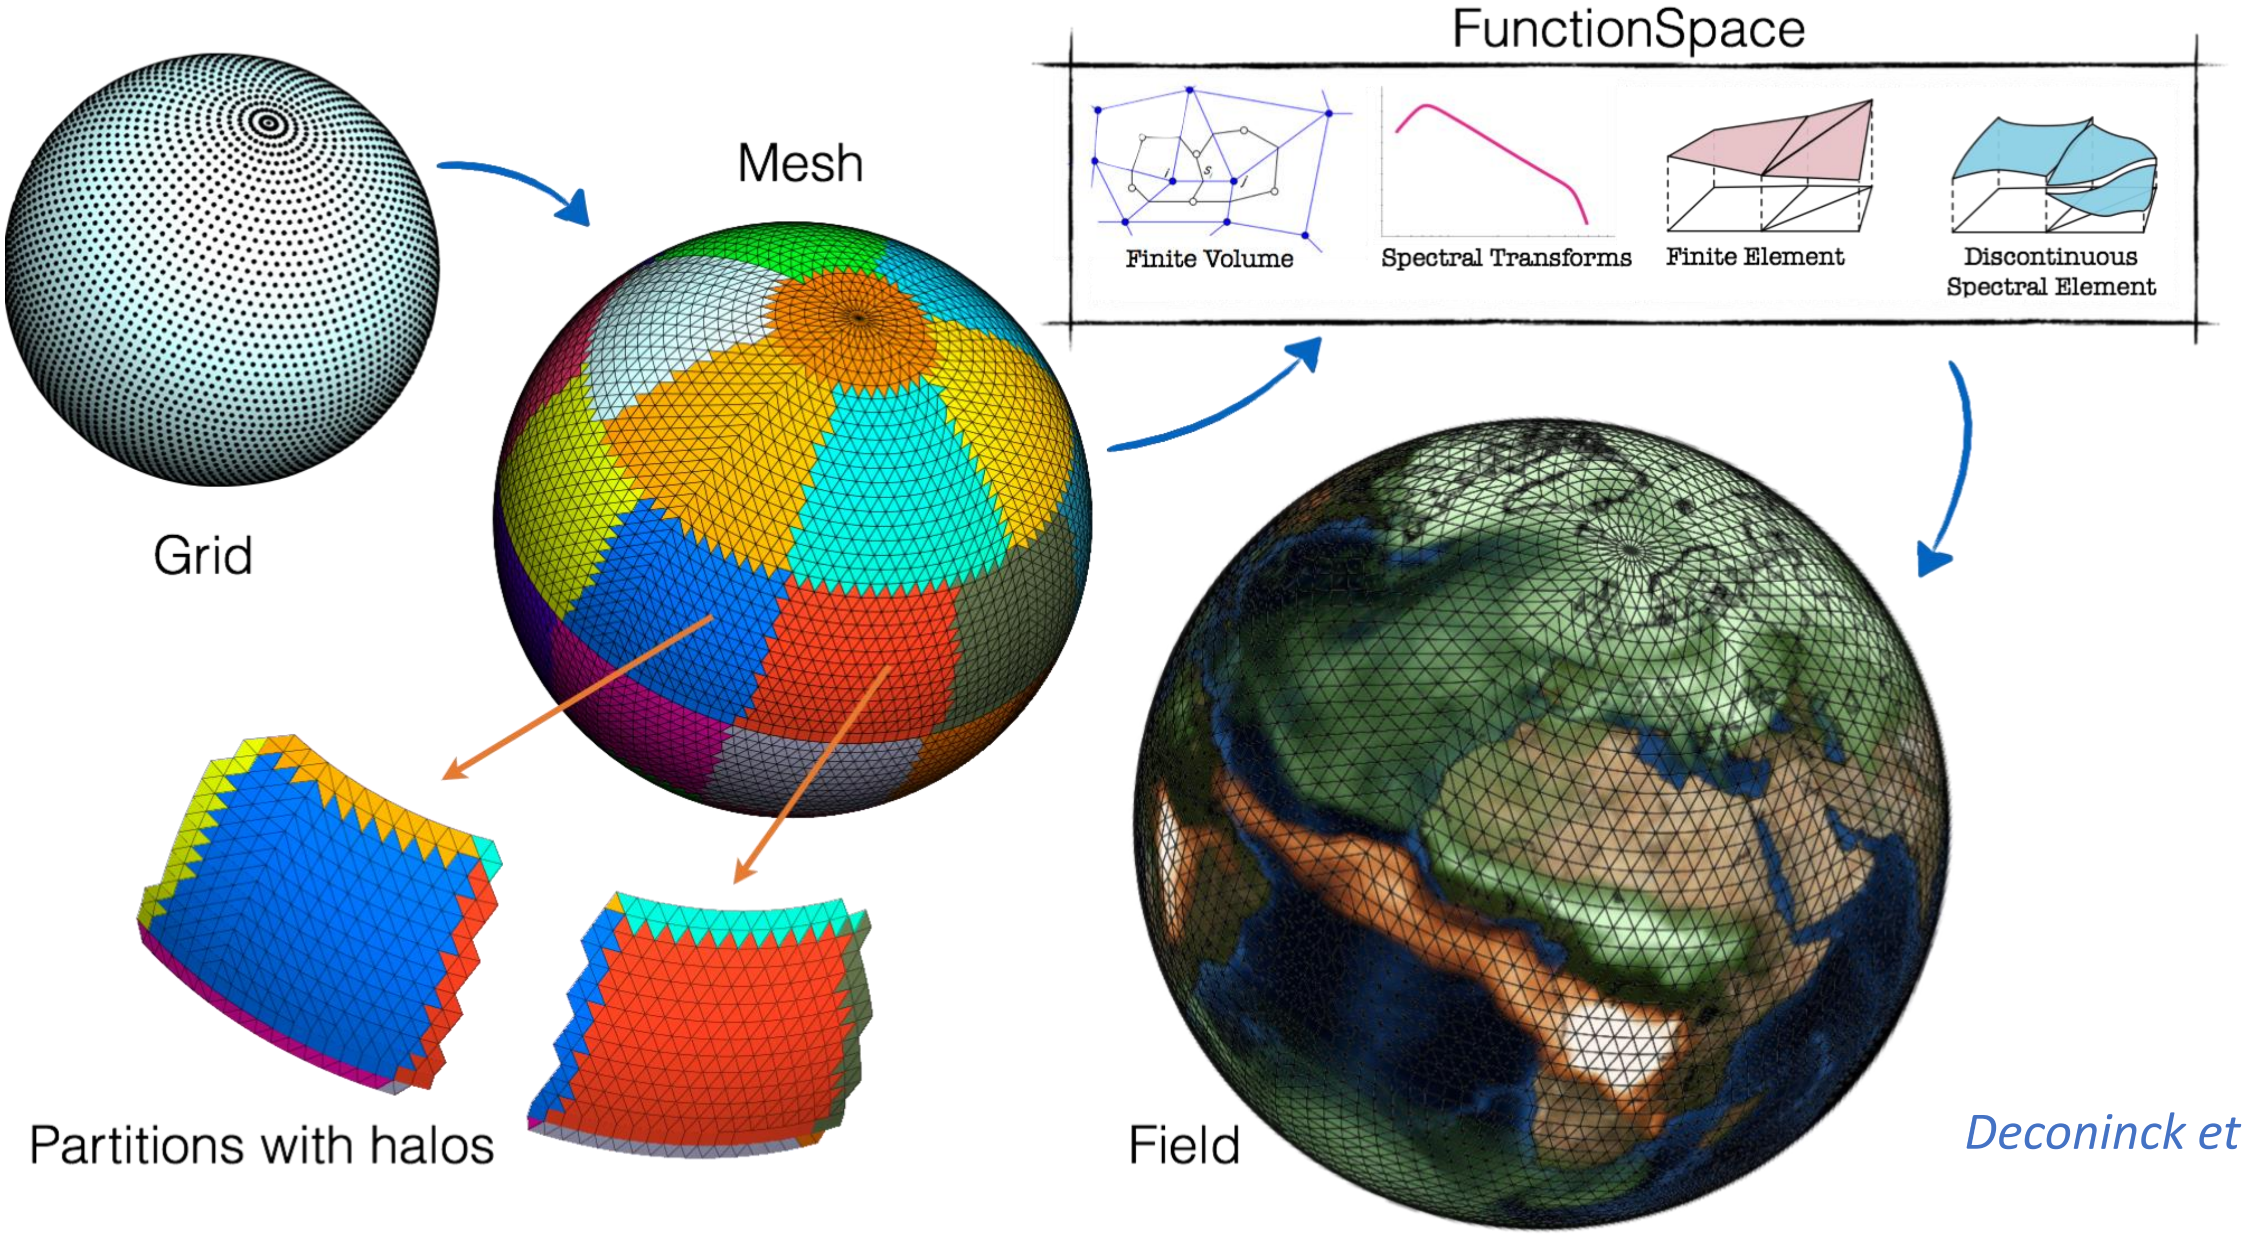

Deconinck et al. 2017

IFS dynamical core options at ECMWF

Christian Kuehnlein

| currently operational |

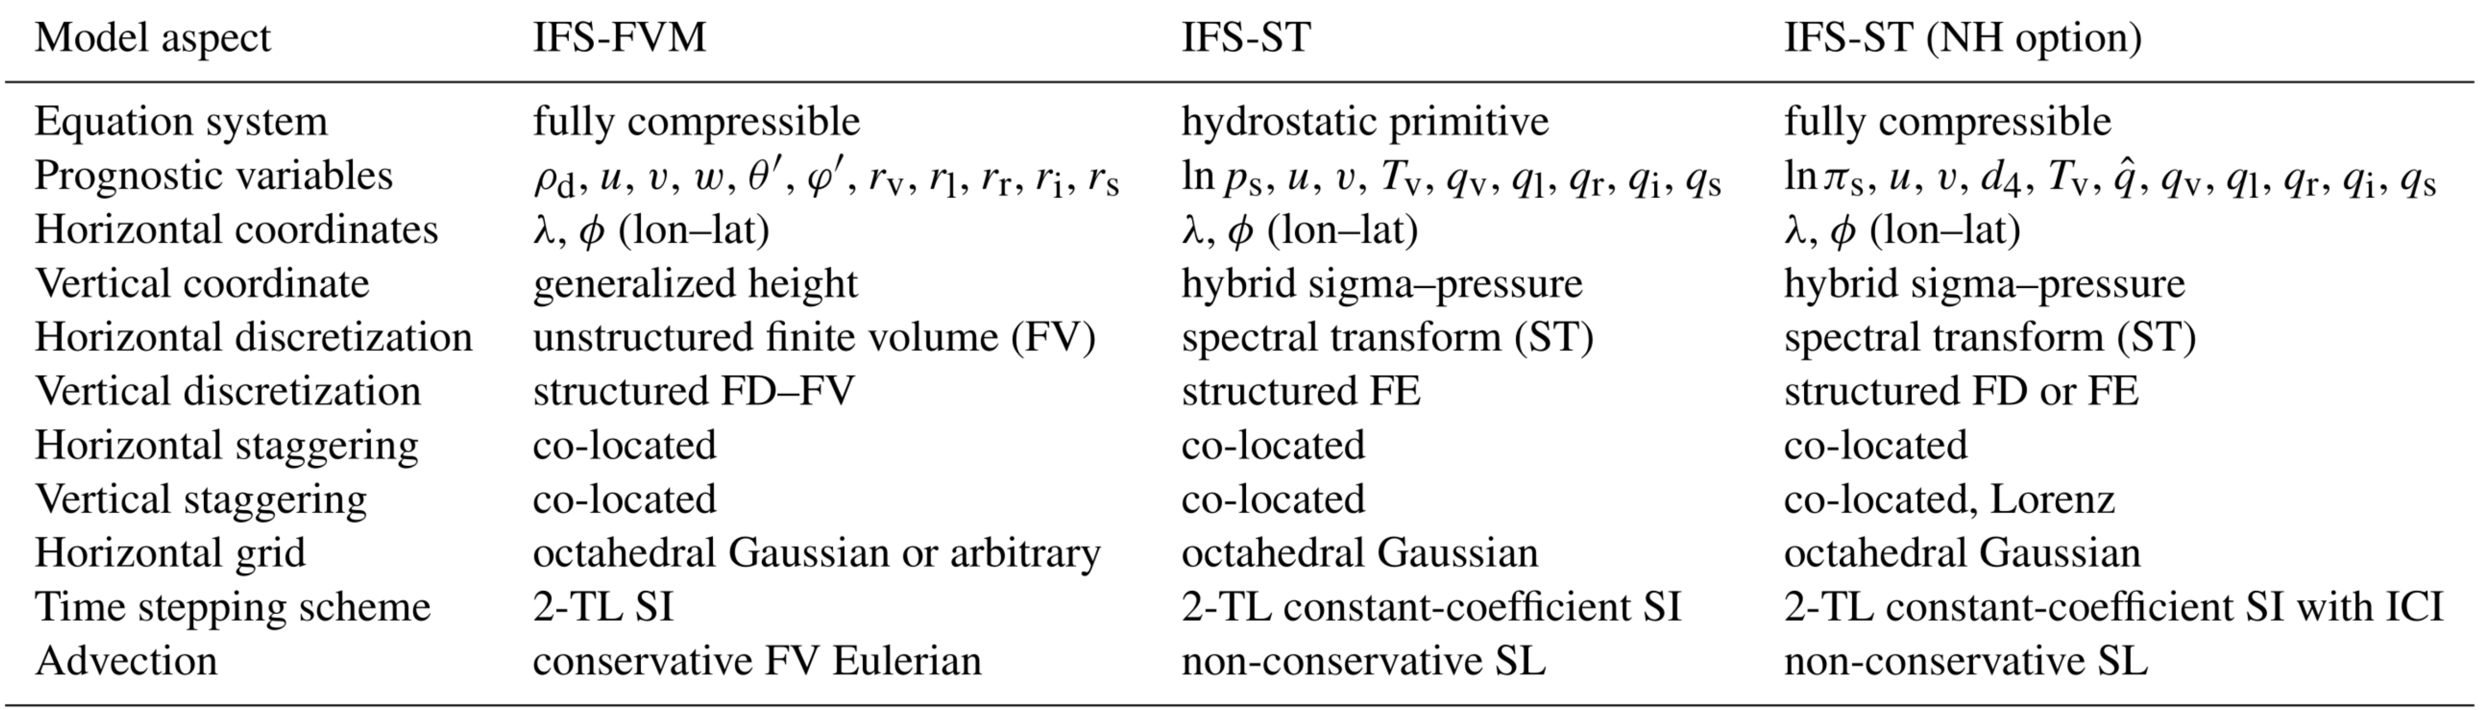

Model aspect	IFS-FVM	IFS-ST	IFS-ST (NH option)
Equation system	fully compressible	hydrostatic primitive	fully compressible
Prognostic variables	$\rho_\mathrm{d}, u, v, w, \theta', \varphi', r_\mathrm{v}, r_\mathrm{l}, r_\mathrm{r}, r_\mathrm{i}, r_\mathrm{s}$	$\ln p_\mathrm{s}, u, v, T_\mathrm{v}, q_\mathrm{v}, q_\mathrm{l}, q_\mathrm{r}, q_\mathrm{i}, q_\mathrm{s}$	$\ln \pi_\mathrm{s}, u, v, d_4, T_\mathrm{v}, \hat{q}, q_\mathrm{v}, q_\mathrm{l}, q_\mathrm{r}, q_\mathrm{i}, q_\mathrm{s}$
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

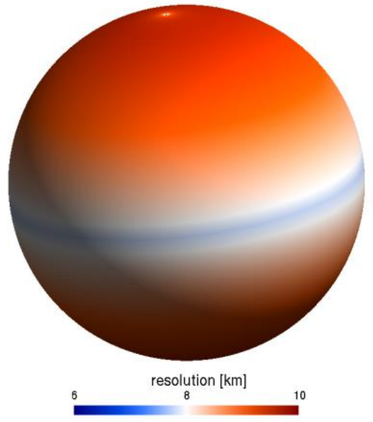

Atlas mesh generation and flexibility through alternative function spaces

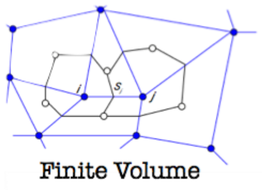

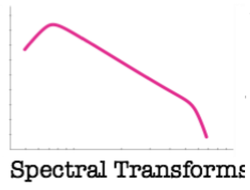

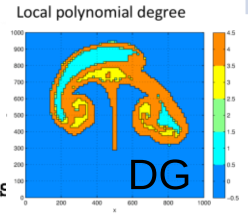

The next 10 years …

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

Michail Diamantakis

Performance and portability

M. Lange, O. Marsden, B. Reuter

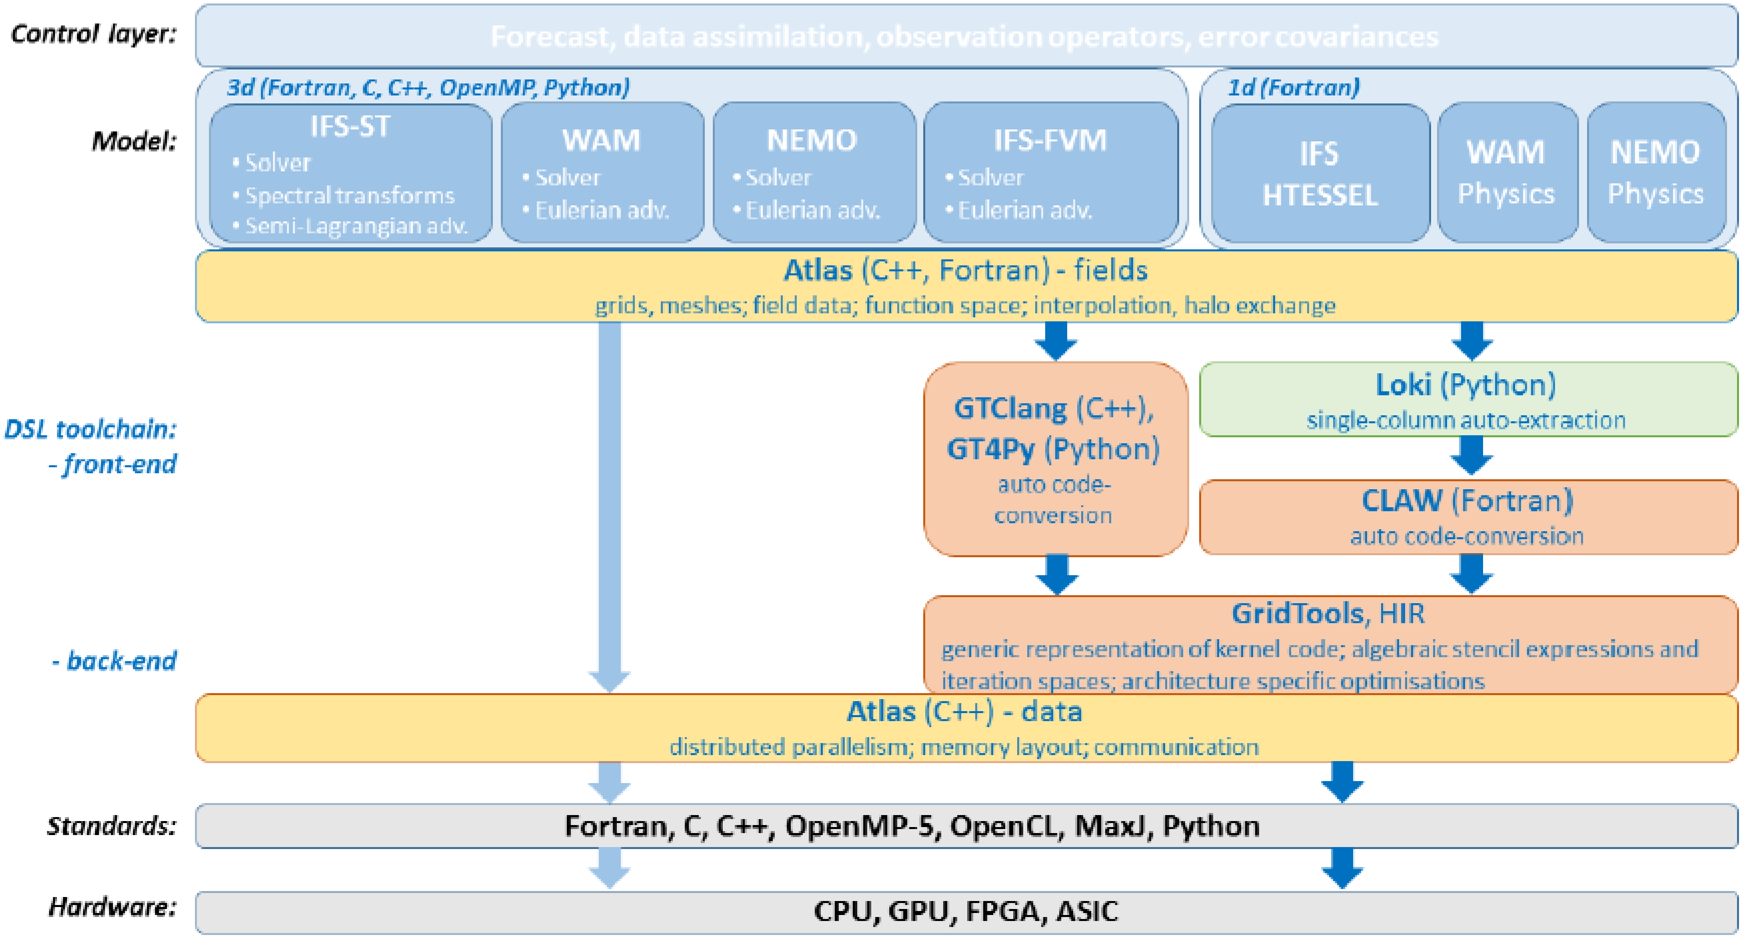

Control layer: Forecast, data assimilation, observation operators, error covariances

Model:

3d (Fortran, C, C++, OpenMP, Python)

IFS-ST
- Solver
- Spectral transforms
- Semi-Lagrangian adv.

WAM
- Solver
- Eulerian adv.

NEMO
- Solver
- Eulerian adv.

IFS-FVM
- Solver
- Eulerian adv.

1d (Fortran)

IFS HTESSEL

WAM Physics

NEMO Physics

Atlas (C++, Fortran) - **fields**
grids, meshes; field data; function space; interpolation, halo exchange

DSL toolchain:
- front-end

GTClang (C++), **GT4Py** (Python)
auto code-conversion

Loki (Python)
single-column auto-extraction

CLAW (Fortran)
auto code-conversion

- back-end

GridTools, HIR
generic representation of kernel code; algebraic stencil expressions and iteration spaces; architecture specific optimisations

Atlas (C++) - **data**
distributed parallelism; memory layout; communication

Standards: Fortran, C, C++, OpenMP-5, OpenCL, MaxJ, Python

Hardware: CPU, GPU, FPGA, ASIC

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

ECMWF Scalability Programme

nextgenio

LEXIS

ESCAPE

EuroEXA

MAESTRO DATA ORCHESTRATION

EPiGRAM HS

ESCAPE 2

esiwace2 CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

esiwace CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

ExtremeEarth

DIGITAL TWINS

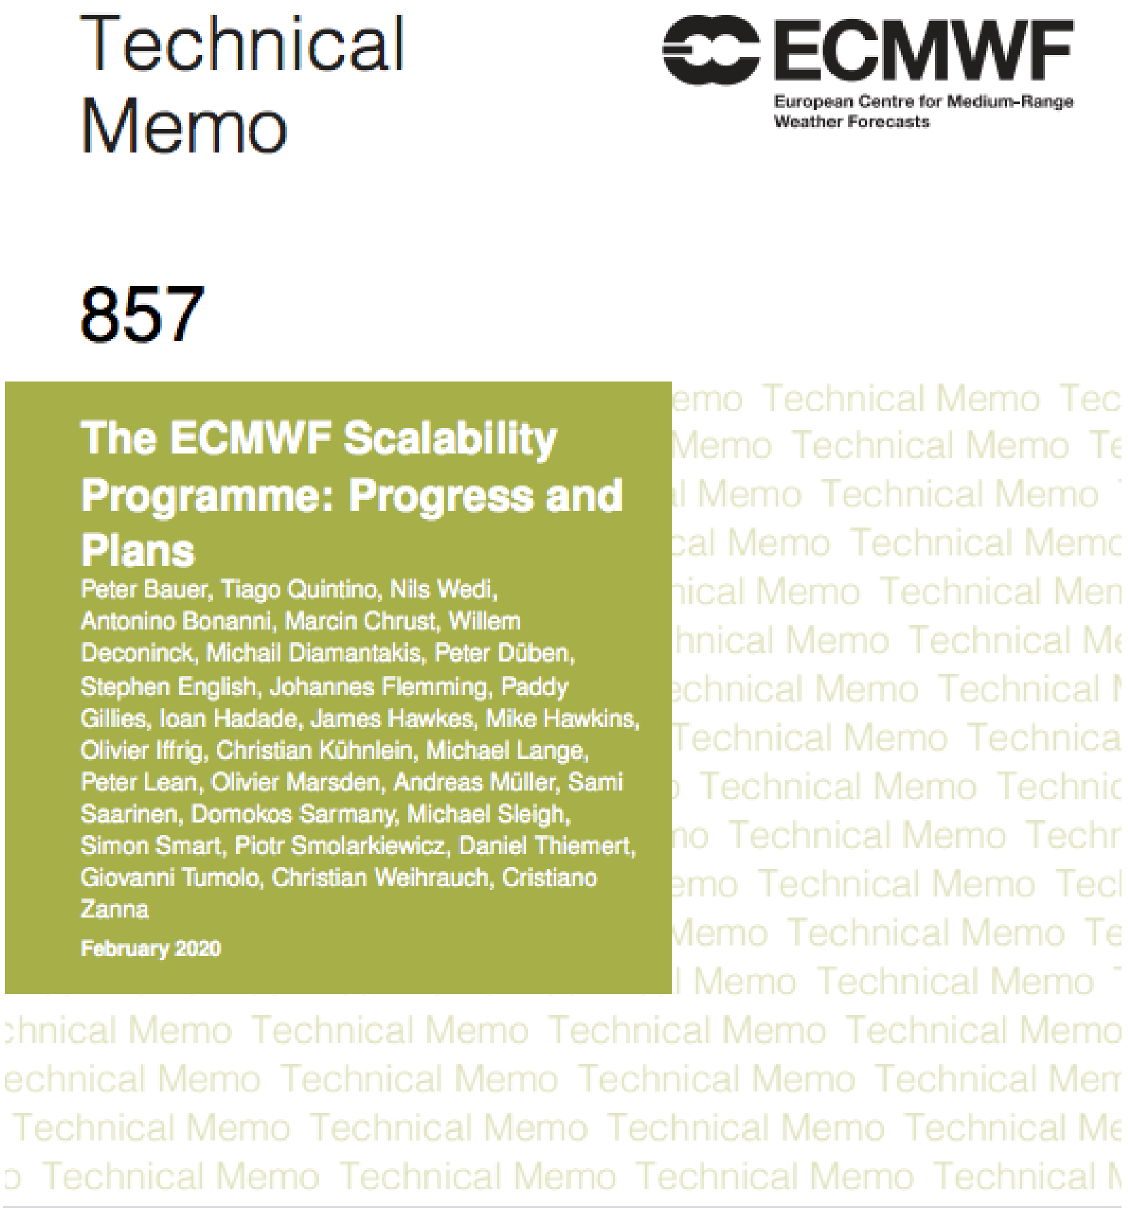

Technical Memo

ECMWF
European Centre for Medium-Range Weather Forecasts

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

46R1 vs 45R1 and single versus double precision

→ Sam Hatfield's talk

(km)

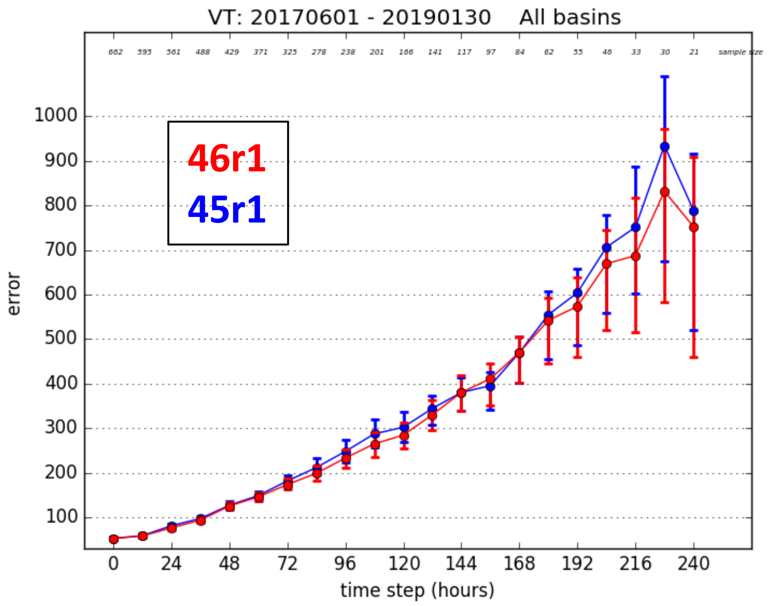

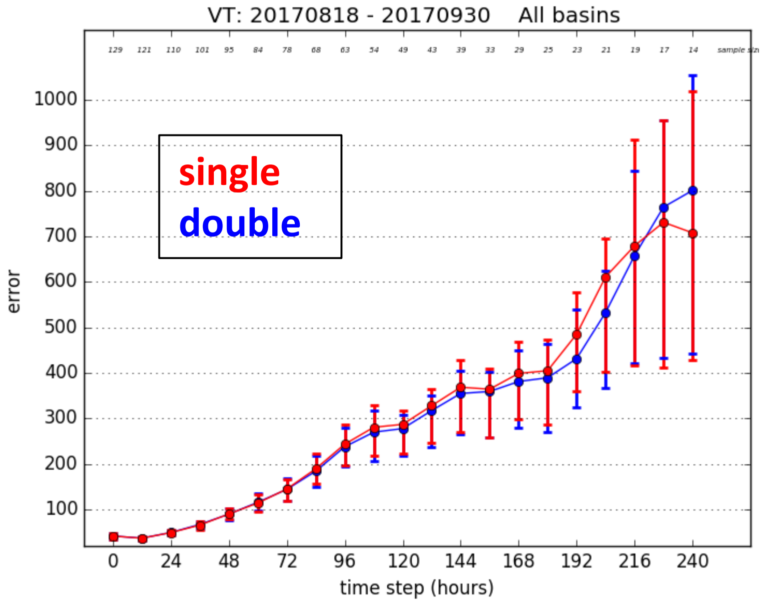

mean position error

(hPa)

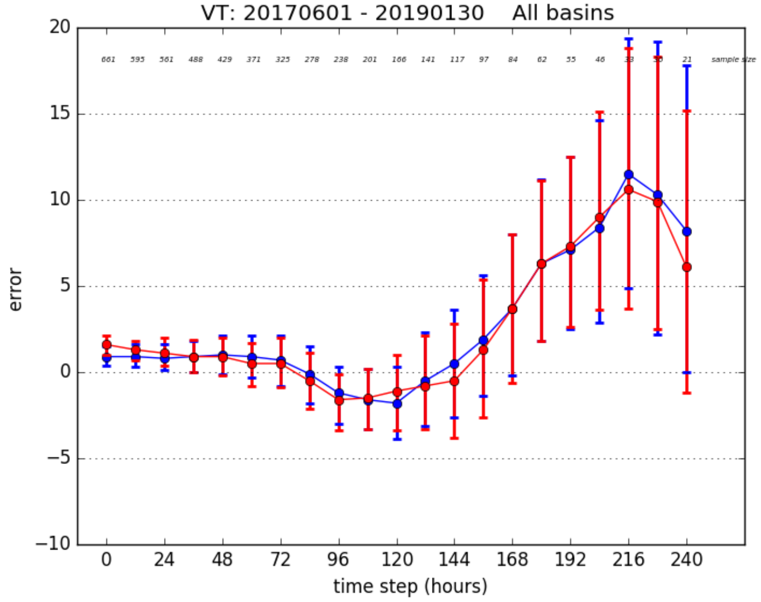

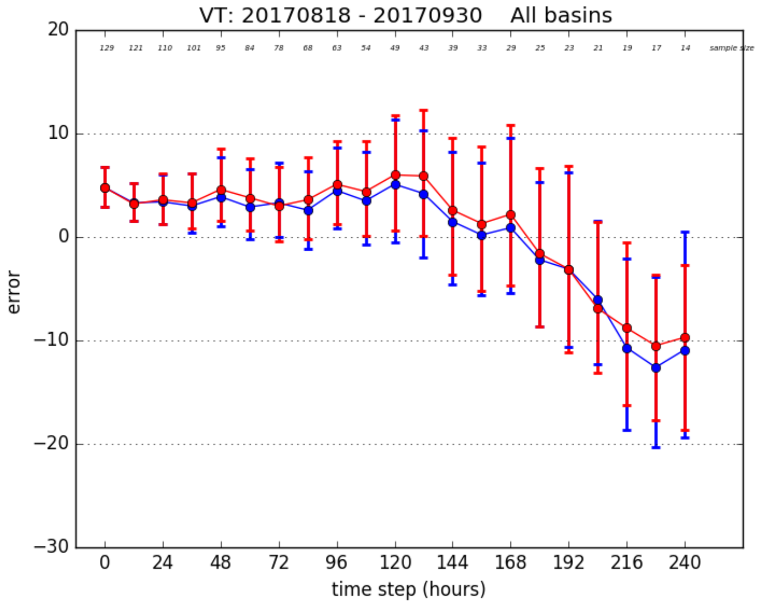

mean intensity error

Vana, Dueben et al, 2017

GPU Spectral transform dwarf

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

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

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

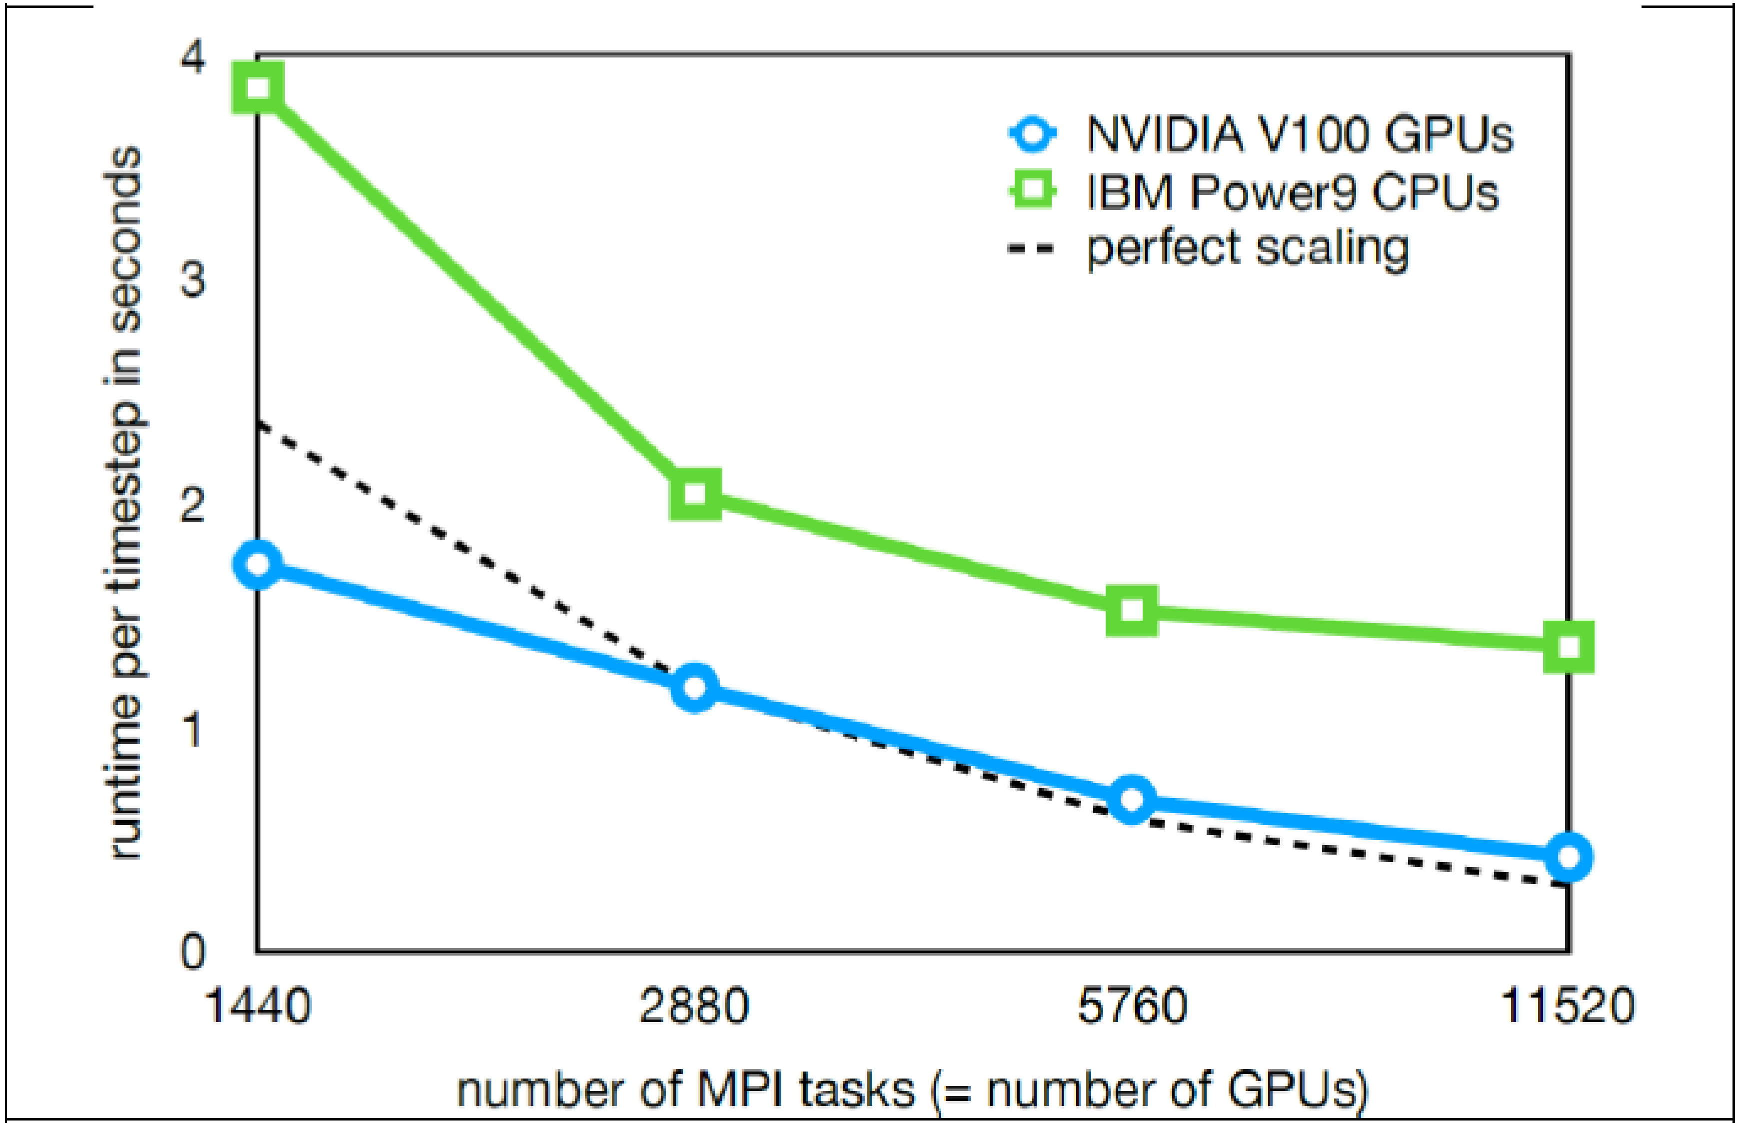

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

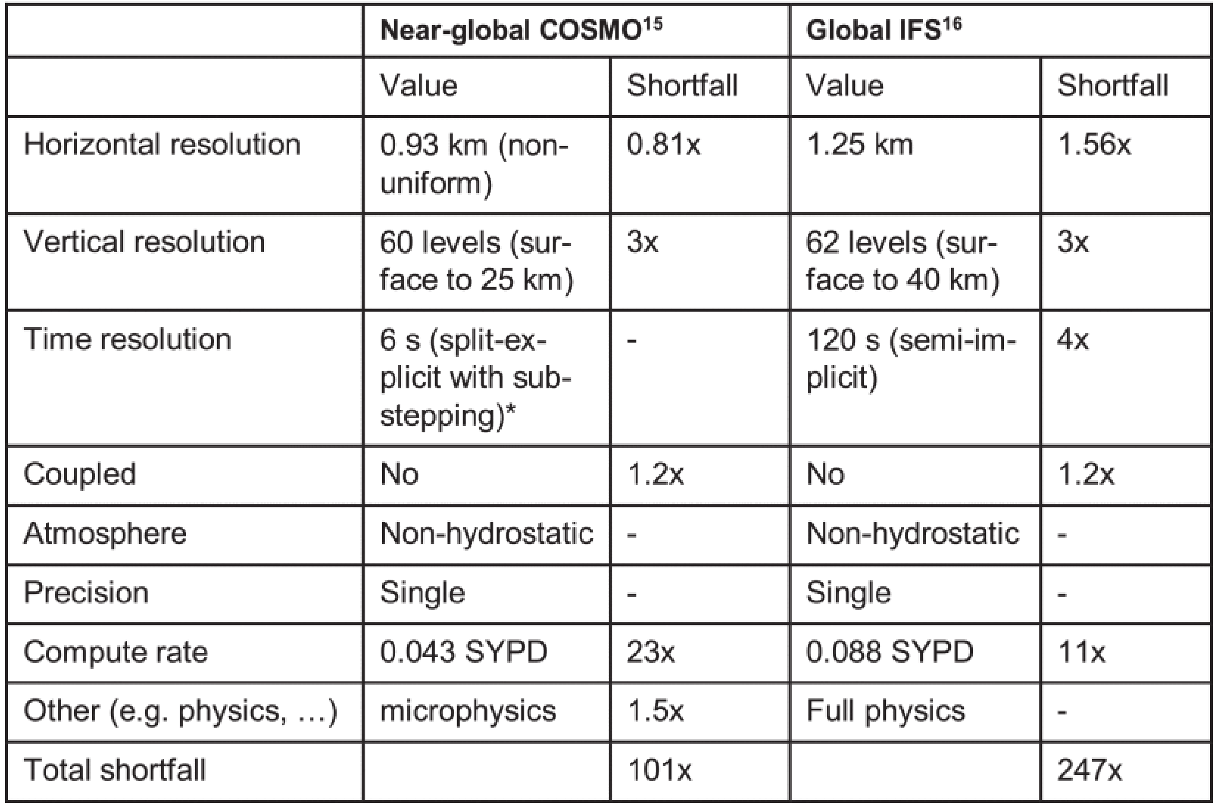

	Near-global COSMO[15]		Global IFS[16]	
	Value	Shortfall	Value	Shortfall
Horizontal resolution	0.93 km (non-uniform)	0.81x	1.25 km	1.56x
Vertical resolution	60 levels (surface to 25 km)	3x	62 levels (surface to 40 km)	3x
Time resolution	6 s (split-explicit with substepping)*	-	120 s (semi-implicit)	4x
Coupled	No	1.2x	No	1.2x
Atmosphere	Non-hydrostatic	-	Non-hydrostatic	-
Precision	Single	-	Single	-
Compute rate	0.043 SYPD	23x	0.088 SYPD	11x
Other (e.g. physics, …)	microphysics	1.5x	Full physics	-
Total shortfall		101x		247x

(Schulthess et al, 2019)

Production run on ~1/5th of Summit at 1.4 km

	Shortfall
1.4 km	2x
1	1.3x
60	2x
No	1.2x
Hydrostatic	-
Single	-
0.03 SYPD	33x
Full physics	-
Includes I/O	206x

(2020)

shortfall

Cost distribution per timestep

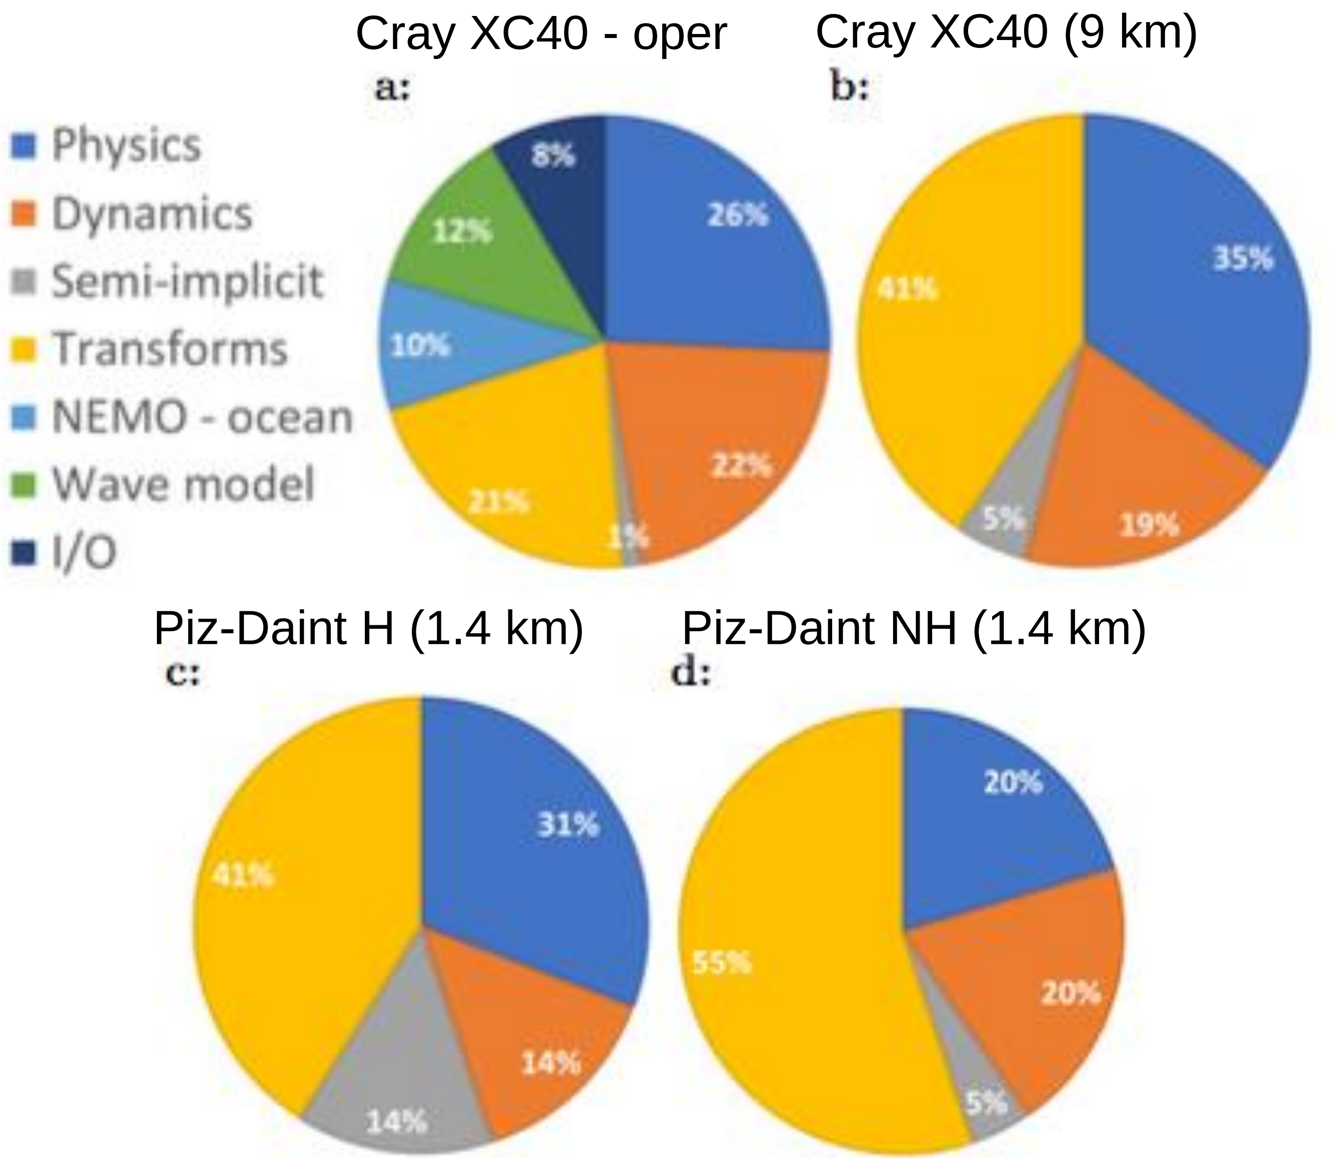
Cray XC40 - oper

Cray XC40 (9 km)

Piz-Daint H (1.4 km)

Piz-Daint NH (1.4 km)

See Sarah-Jane Lock’s talk

Ensemble of assimilations and forecasts

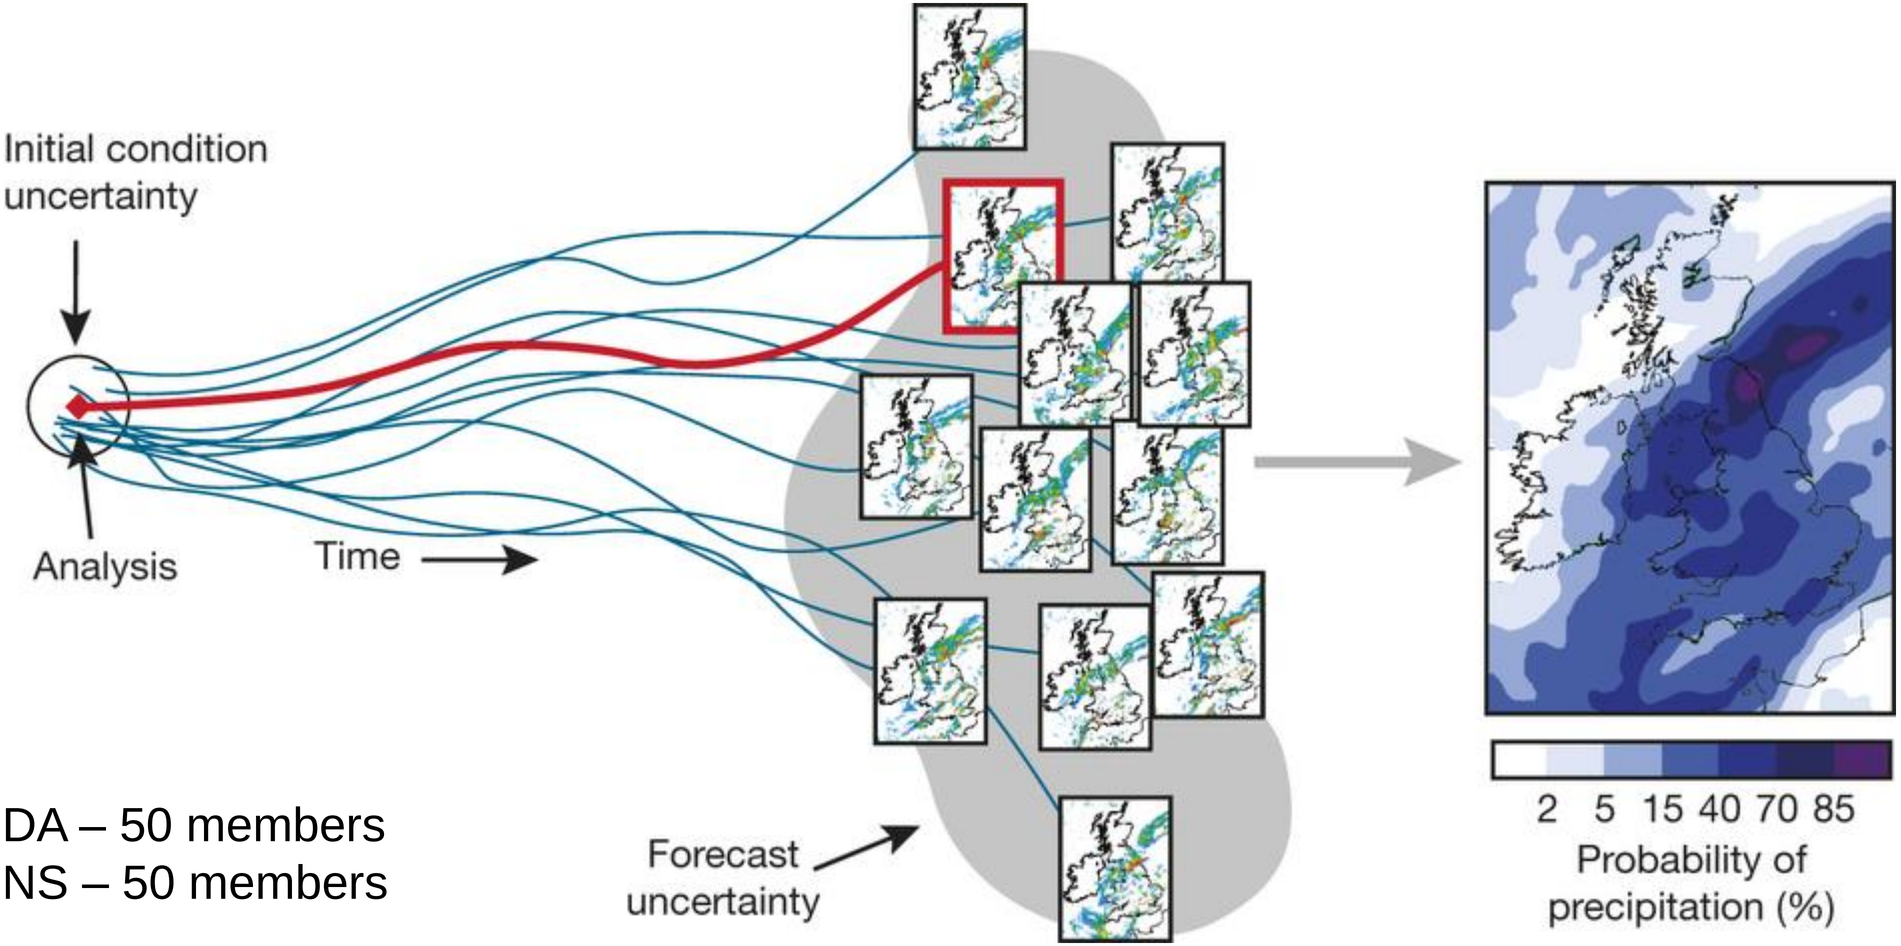

- EDA – 50 members
- ENS – 50 members

Oper, Tco639L91

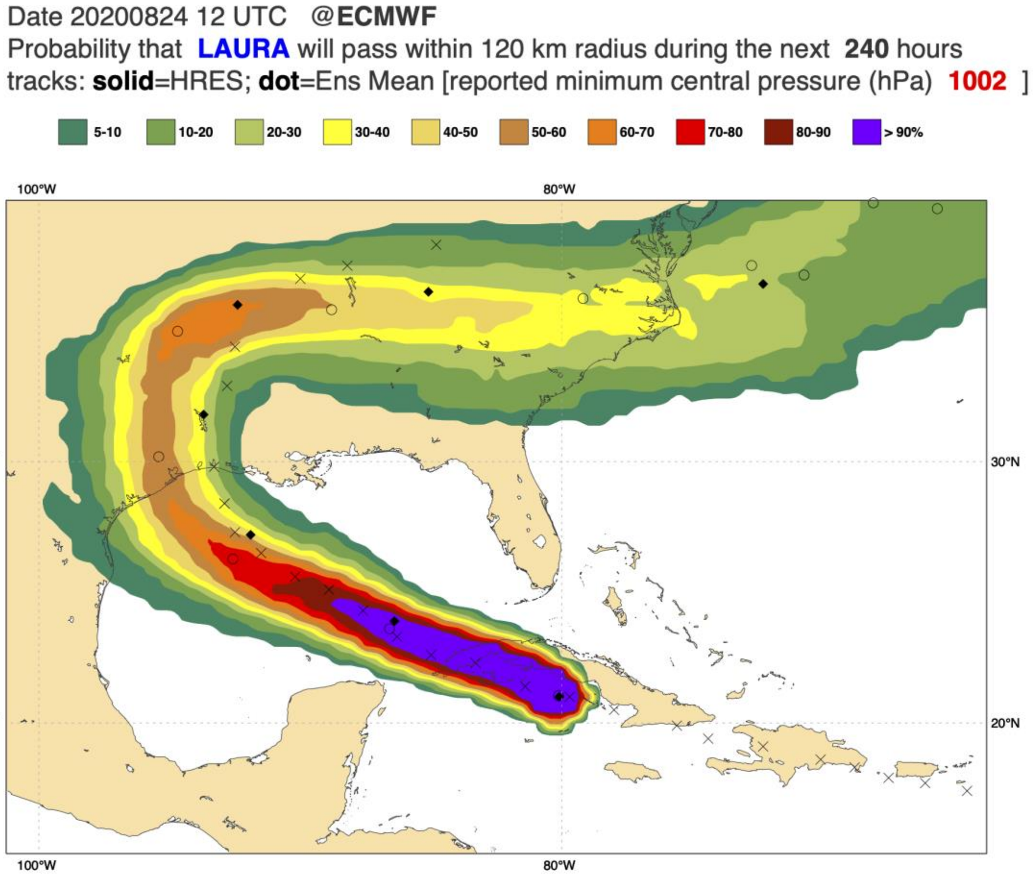

Crosses : observed position
Circle : ensemble mean
Diamonds : Oper HRES

Oper, Tco1279L137

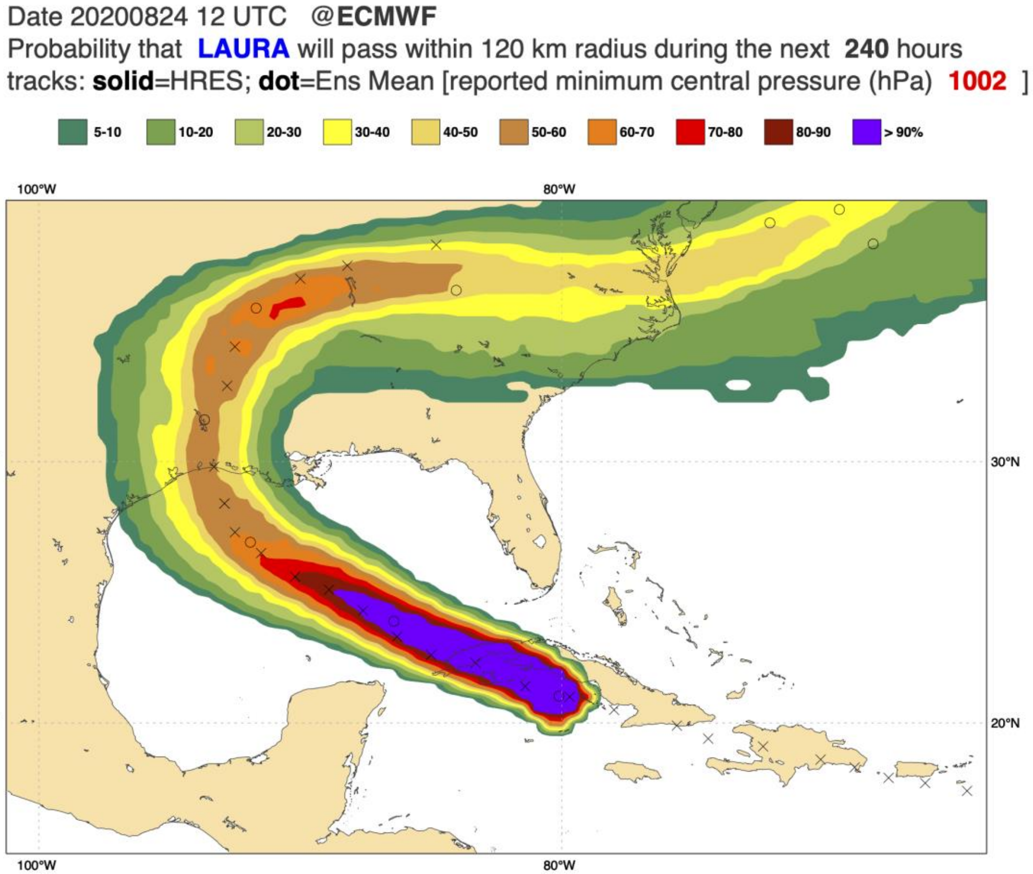

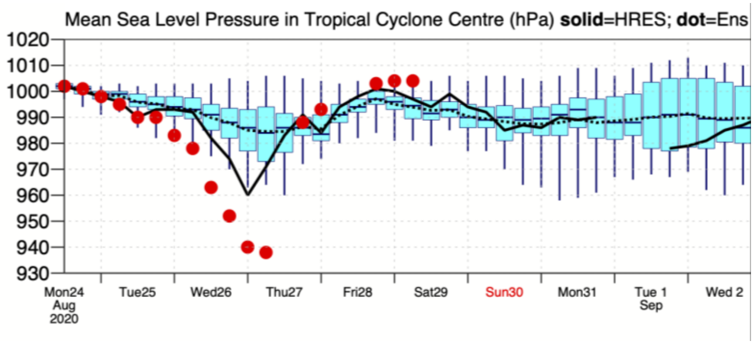

Red dot : observed core pressure
Solid black line : Oper HRES
Box plot : ensemble distribution

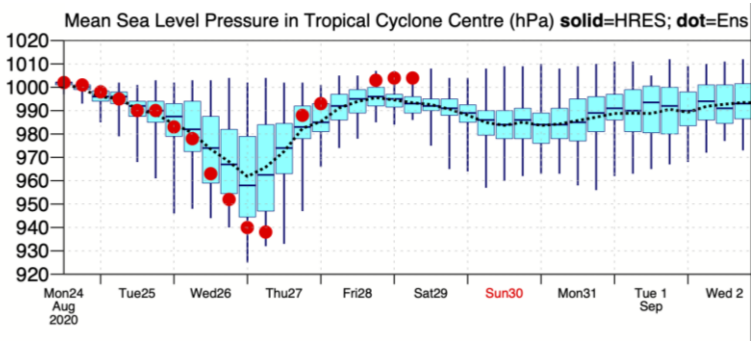

Simon Lang

Ocean – Land – Atmosphere – Sea ice

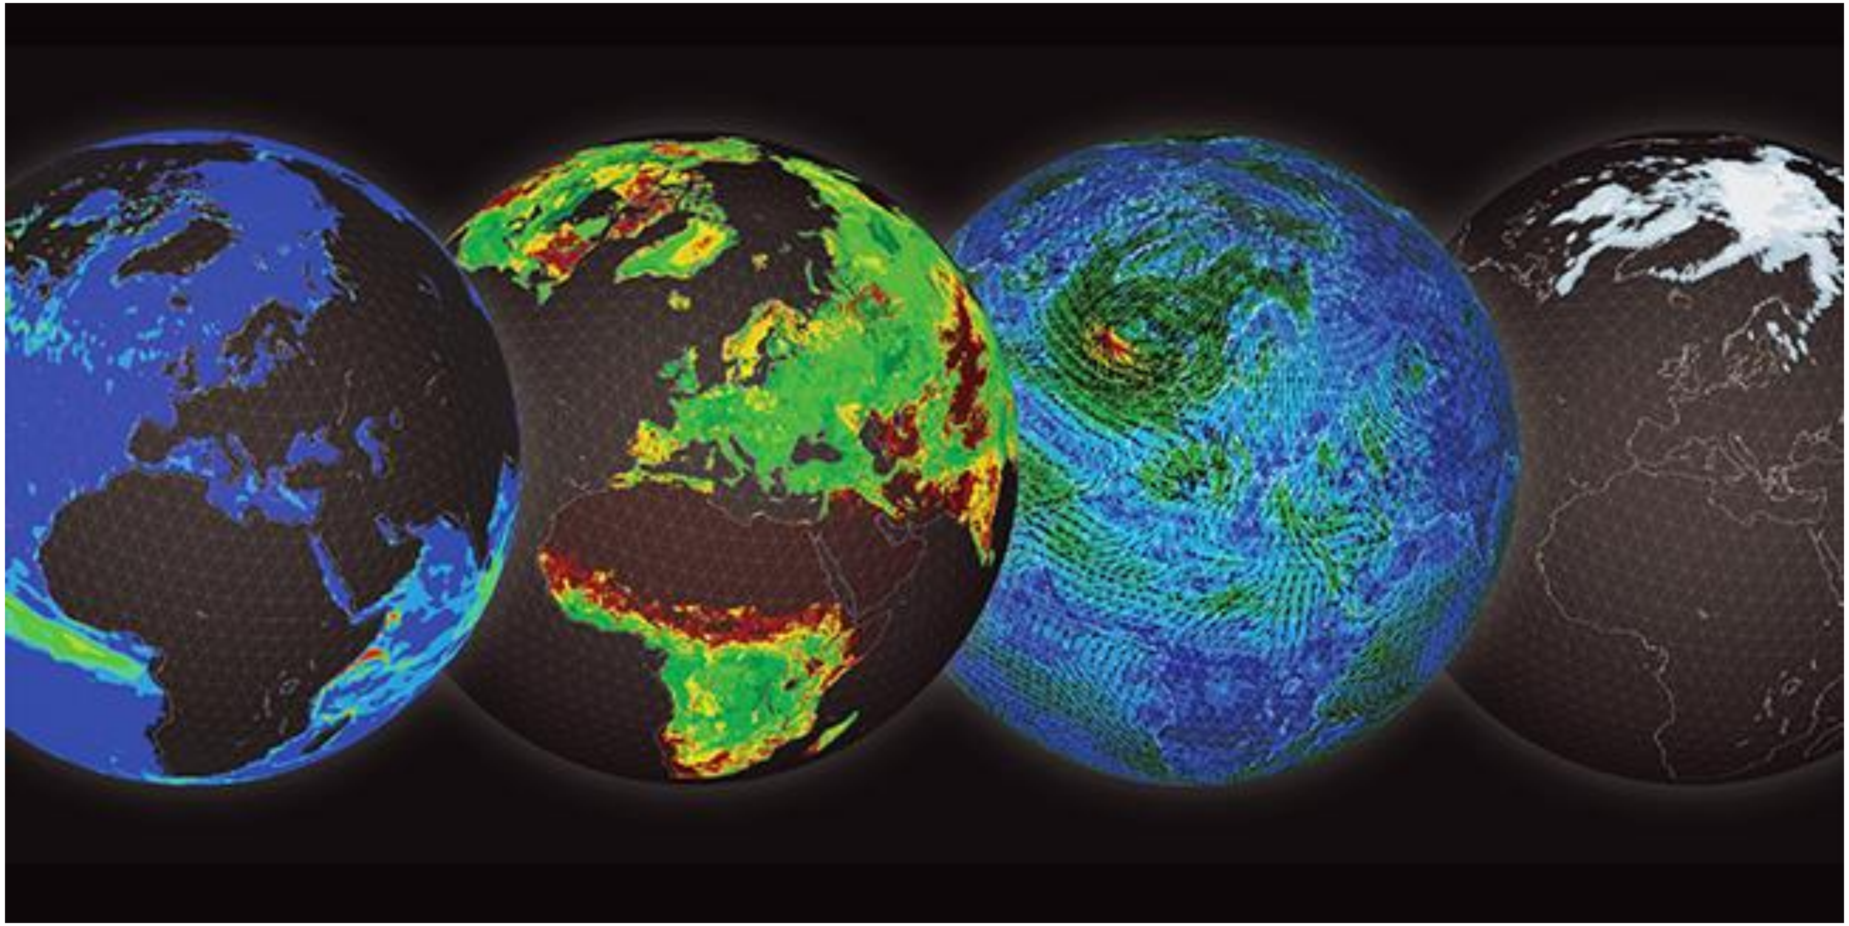

Ocean model - resolution

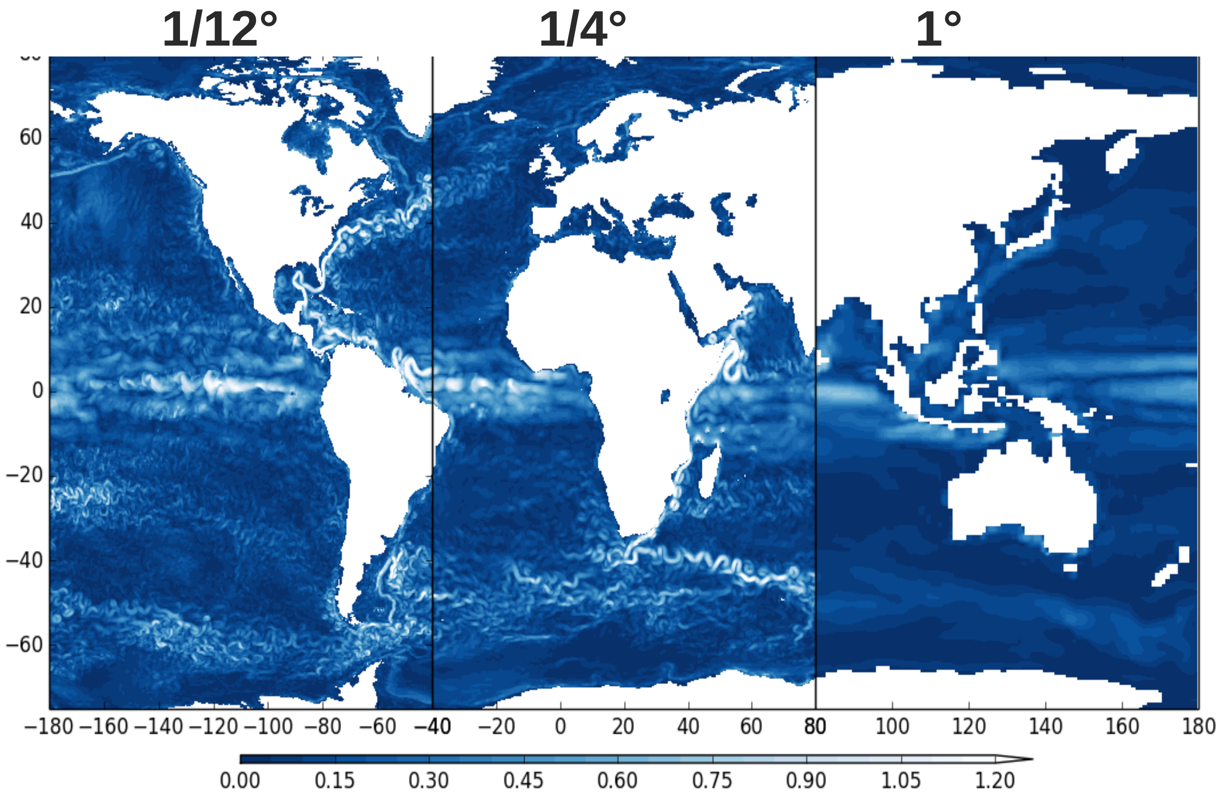

Hewitt et al. (2017)

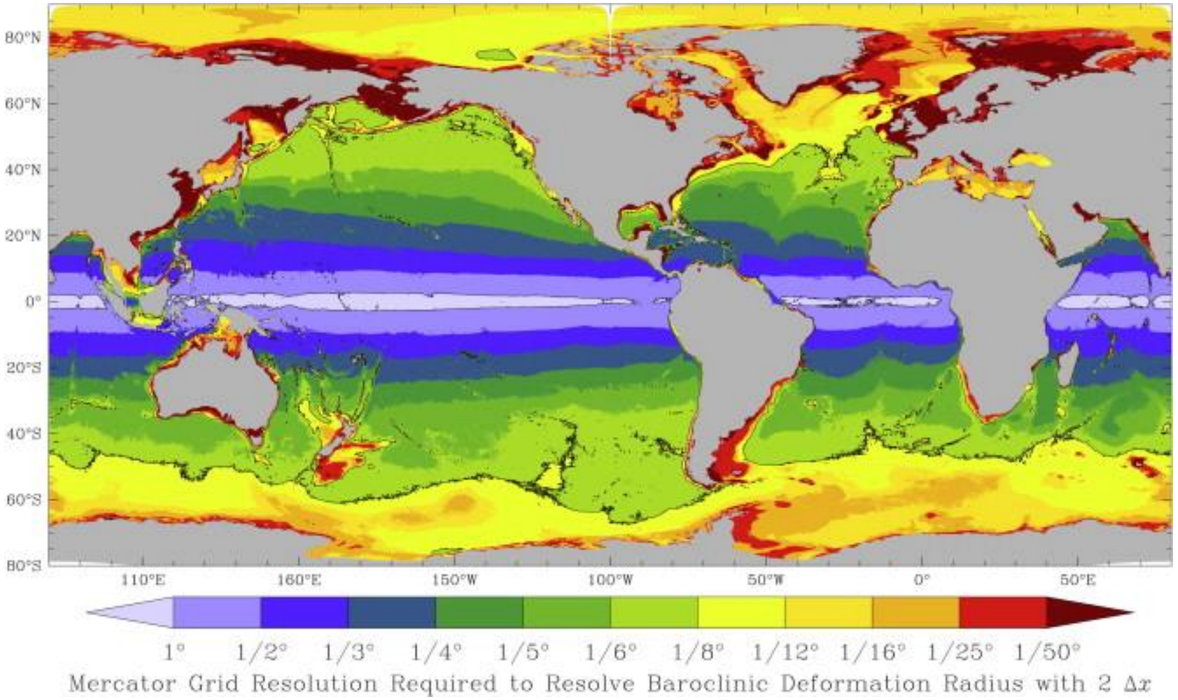

Hallberg (2013)

See talks by S. Danilov, P. Korn, A. Adcroft, J.F. Lemieux…

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

Improved analysis of
- skin temperature
- anthropogenic emissions

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

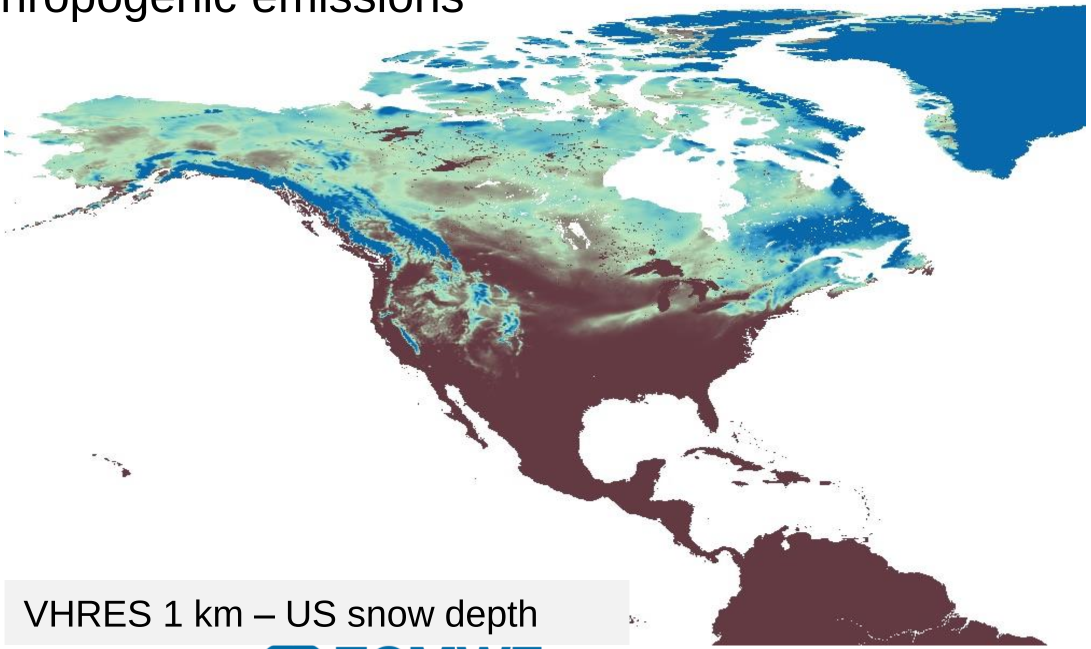

VHRES 1 km – US snow depth

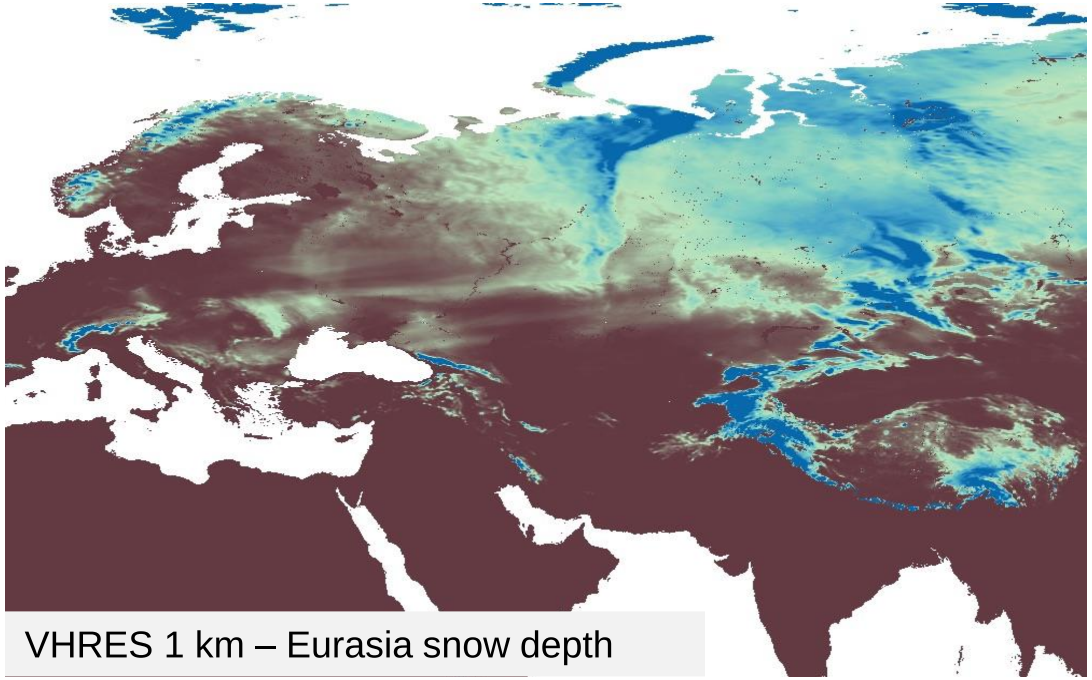

VHRES 1 km – Eurasia snow depth

ECMWF (1km grid-spacing) 48h global forecast

Forecast aerodynamic roughness near surface

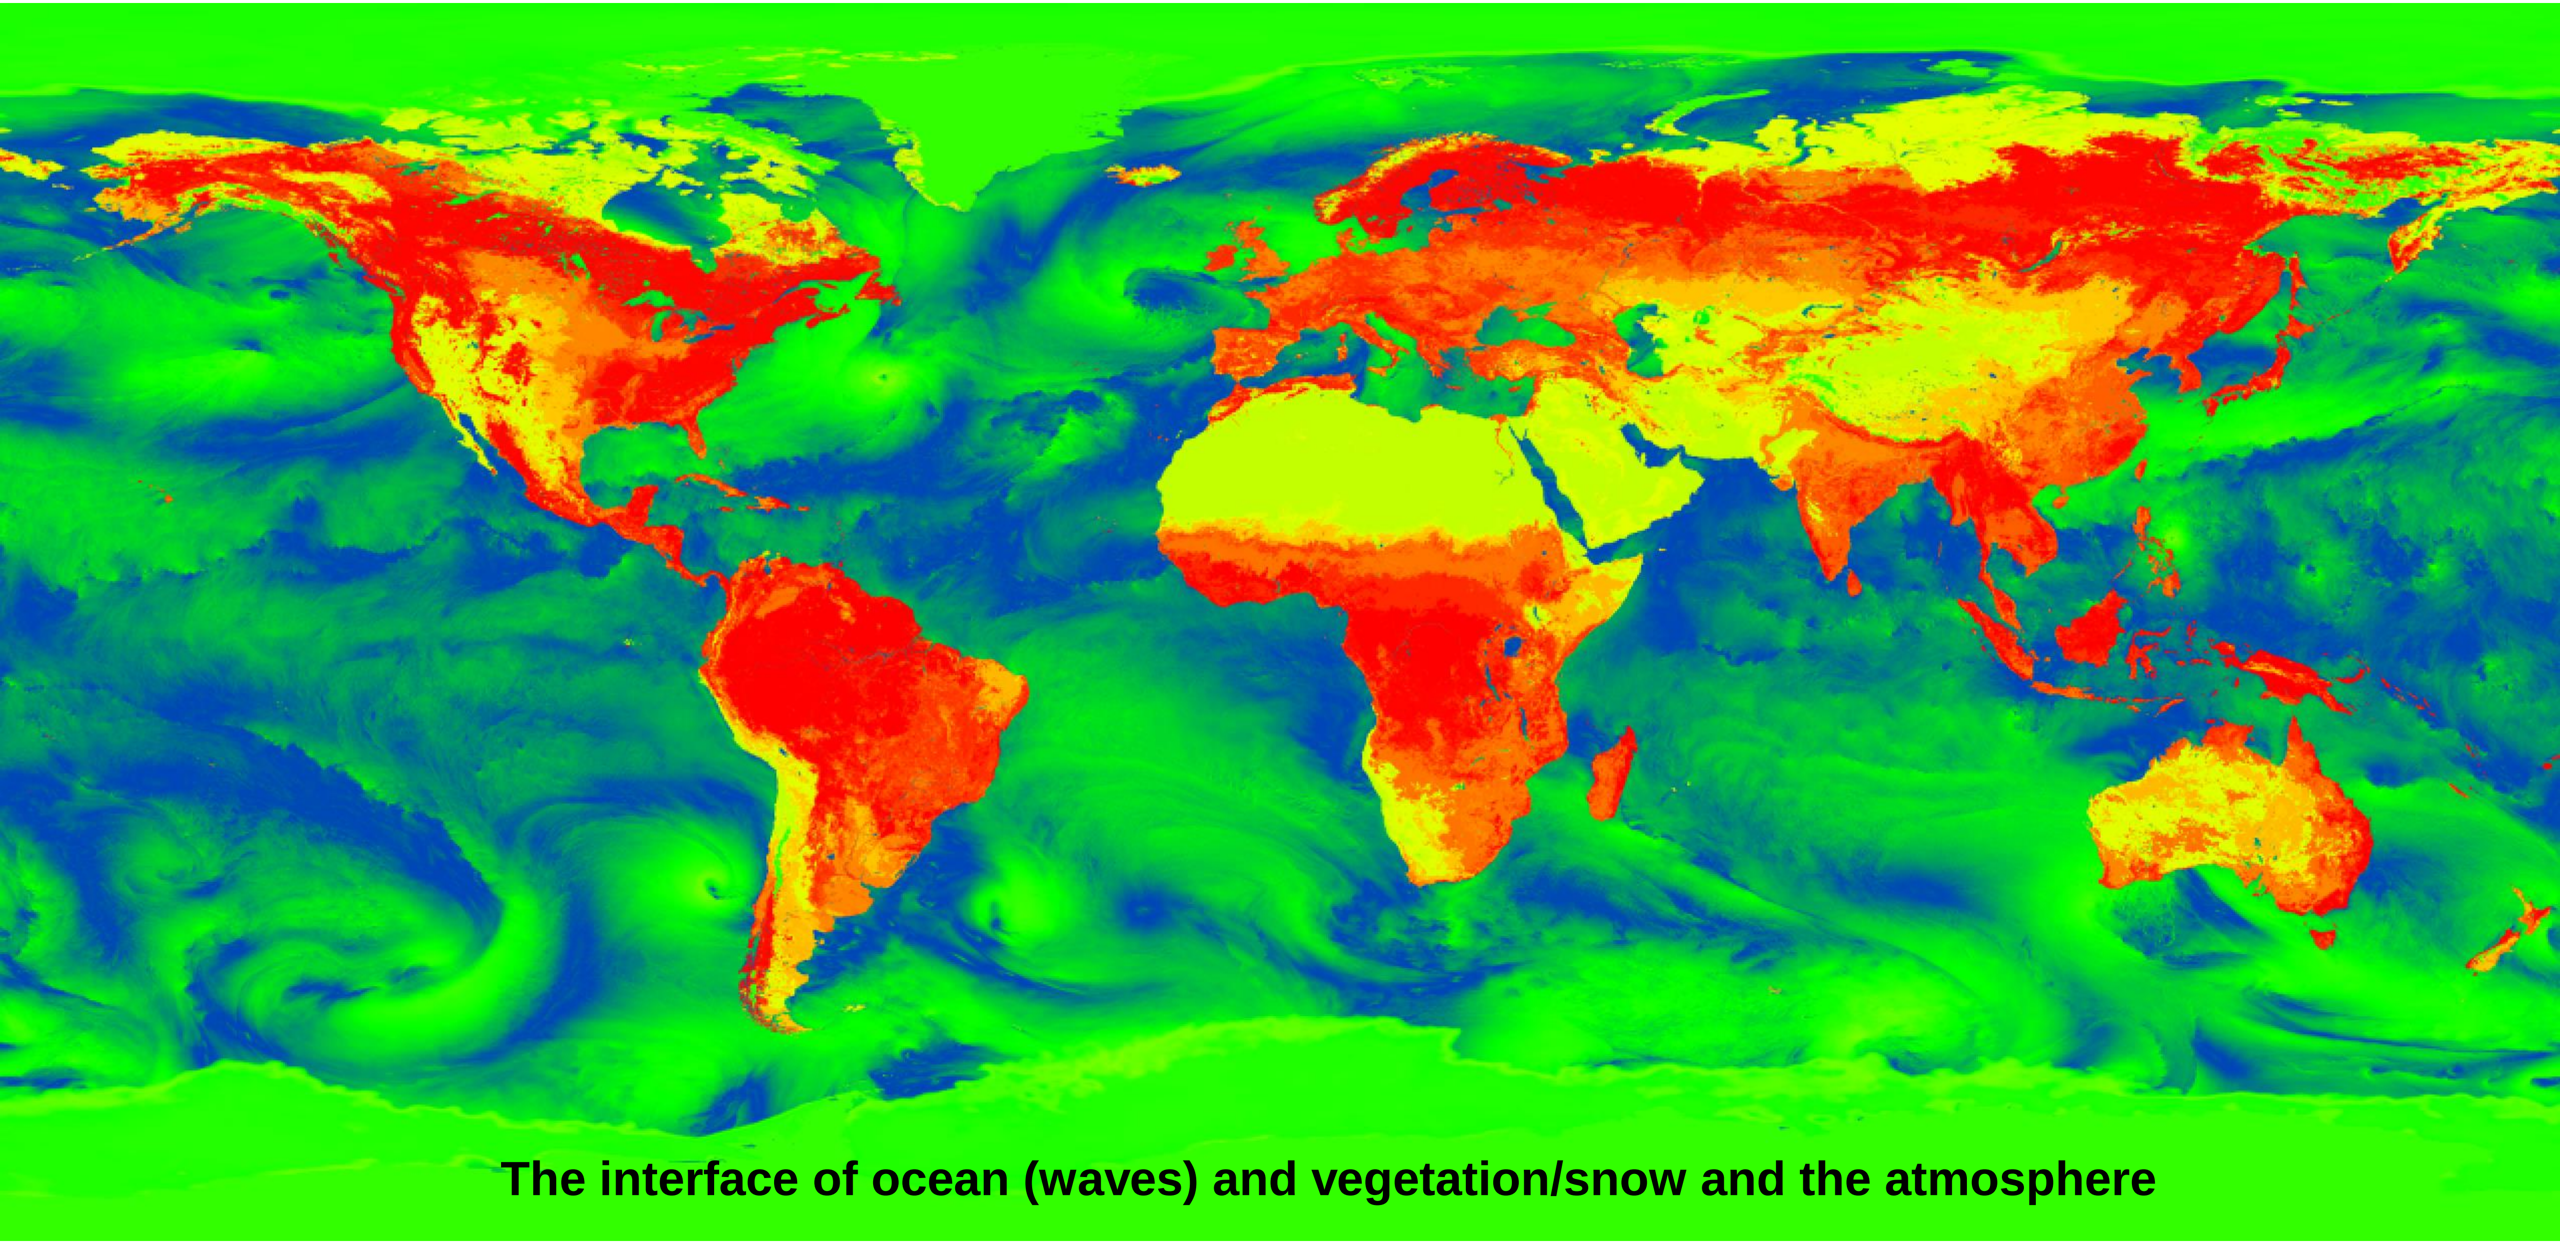

The interface of ocean (waves) and vegetation/snow and the atmosphere

Island wakes

Sea surface temperature

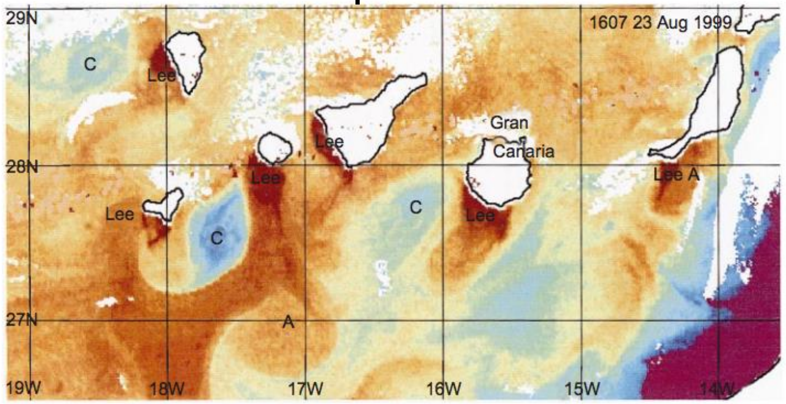

https://digital.csic.es/bitstream/10261/26947/1/0140.pdf

(Strato-)cumulus clouds

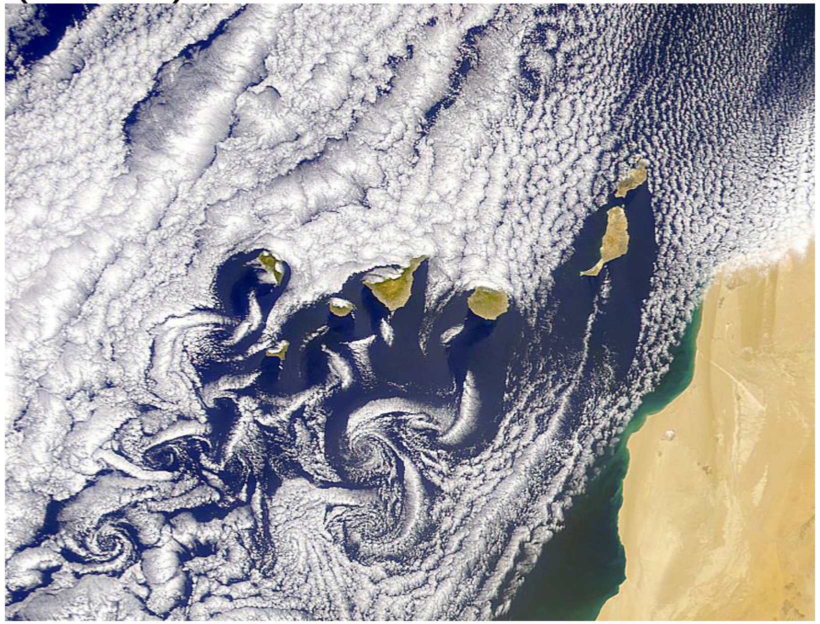

Provided by the SeaWiFS Project, NASA/Goddard Space Flight Center, and ORBIMAGE

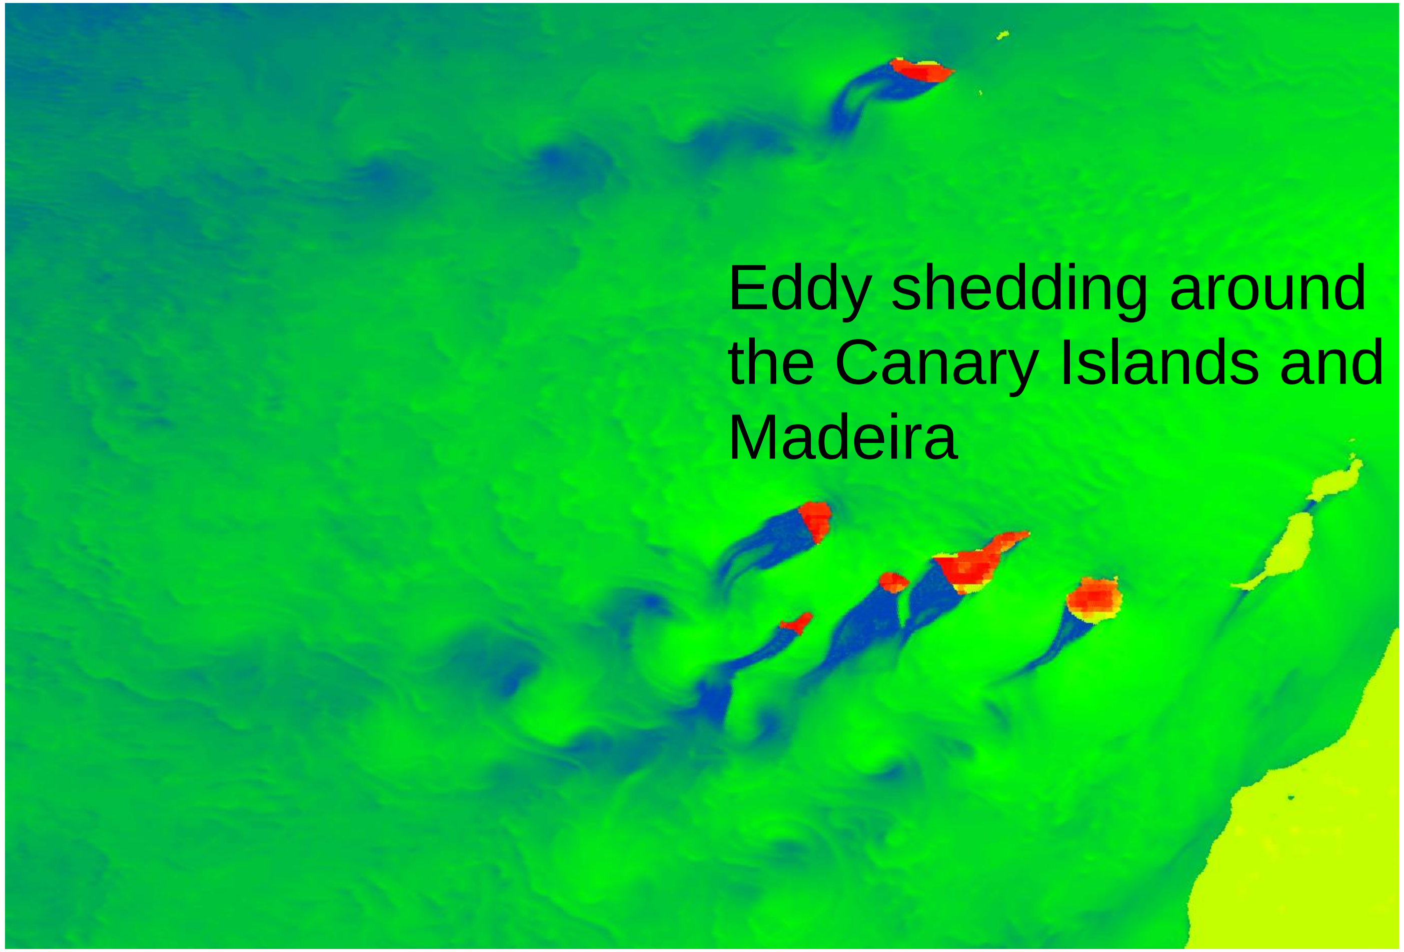

Eddy shedding around the Canary Islands and Madeira

Forecast aerodynamic roughness in 1.4 km simulations

See talk by A. Roland …

A. Agusti-Panareda, G. Balsamo et al

A multi-scale approach: overview

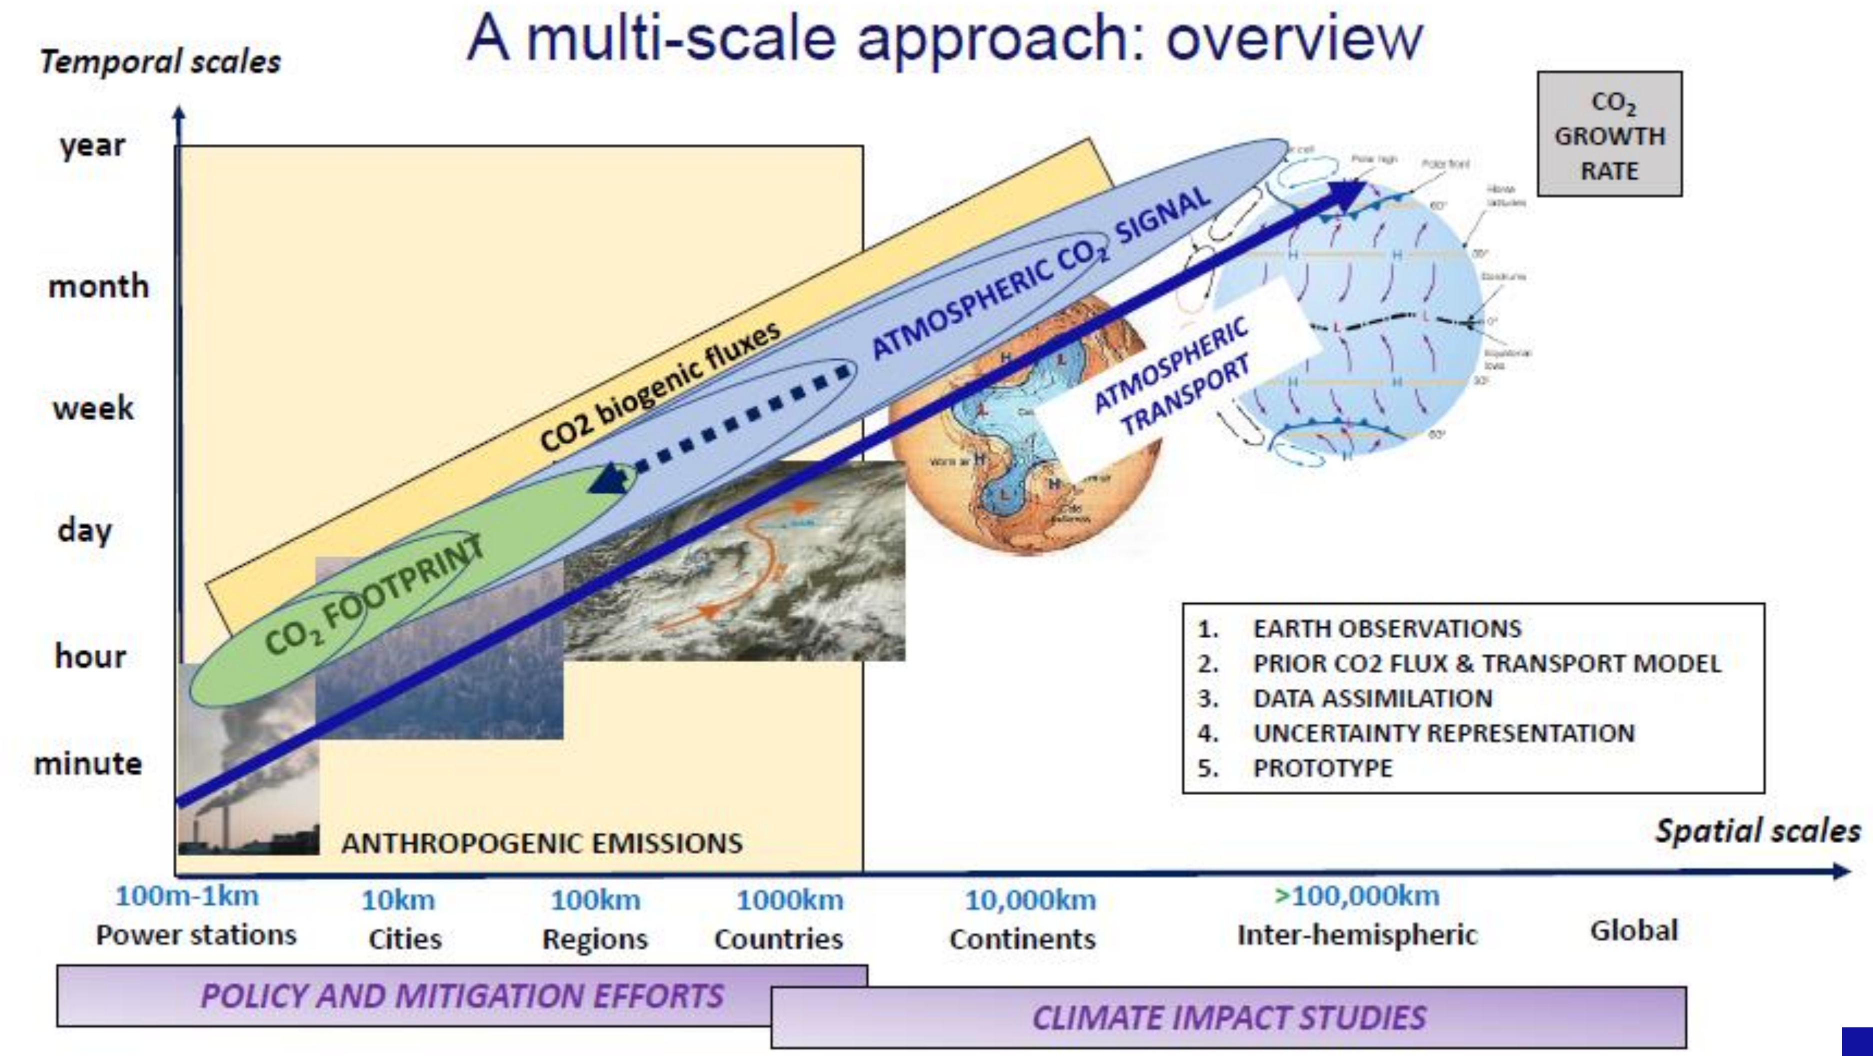

1. EARTH OBSERVATIONS
2. PRIOR CO2 FLUX & TRANSPORT MODEL
3. DATA ASSIMILATION
4. UNCERTAINTY REPRESENTATION
5. PROTOTYPE

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: better reanalysis & monitoring

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

NH winter SSWs

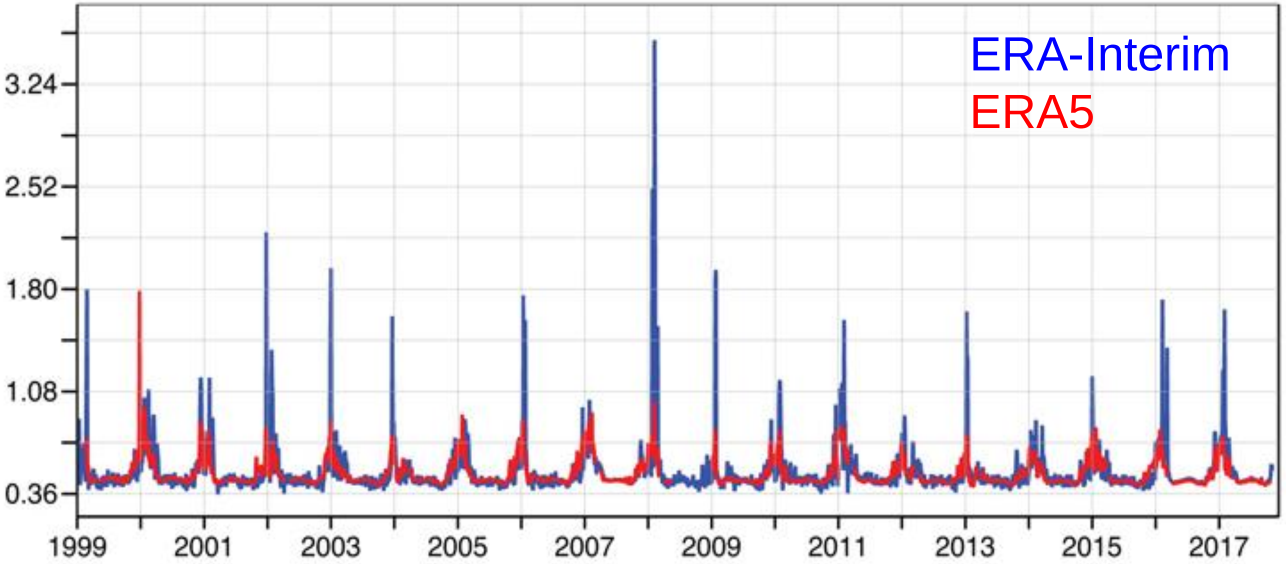

Standard deviation of MW radiances observed vs simulated temperature fields of ERA-Interim (blue) and ERA5 (red) using satellite channel (noaa15) peaking around 5hpa.

T. McNally, A. Simmons

Surface air temperature anomaly for August 2020

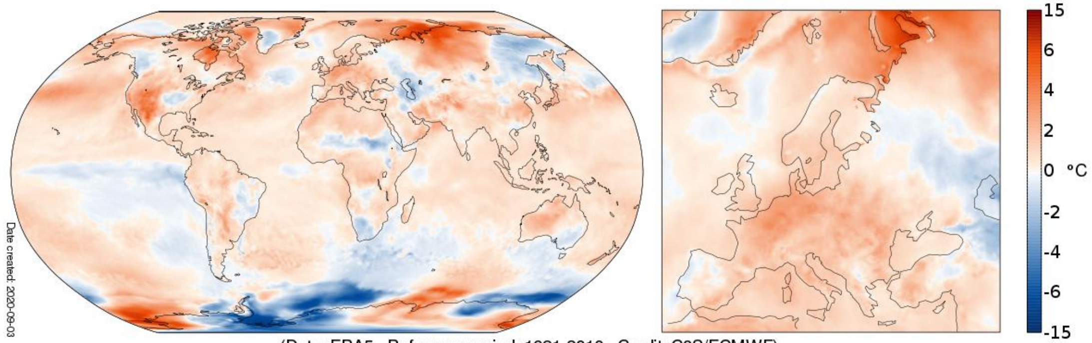

(Data: ERA5. Reference period: 1981-2010. Credit: C3S/ECMWF)

See also Inna Polichtchouk's talk

Break for questions …

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

Four runs to test the impact of time-step-size and hydrostatic equations

Run 1: H, 120s, 0PC; **Run 2:** NH, 120s, 1PC; **Run 3:** H, 30s, 2PC; **Run 4:** NH, 30s, 2PC

Vertical velocity [m/s]

- Run 1 shows the smallest area with strong convection
- Run 1 and 2 are similar
- Run 3 and 4 are similar

-> small differences between H and NH

-> large differences as Function of time step choice

East Africa

South America

Indonesia

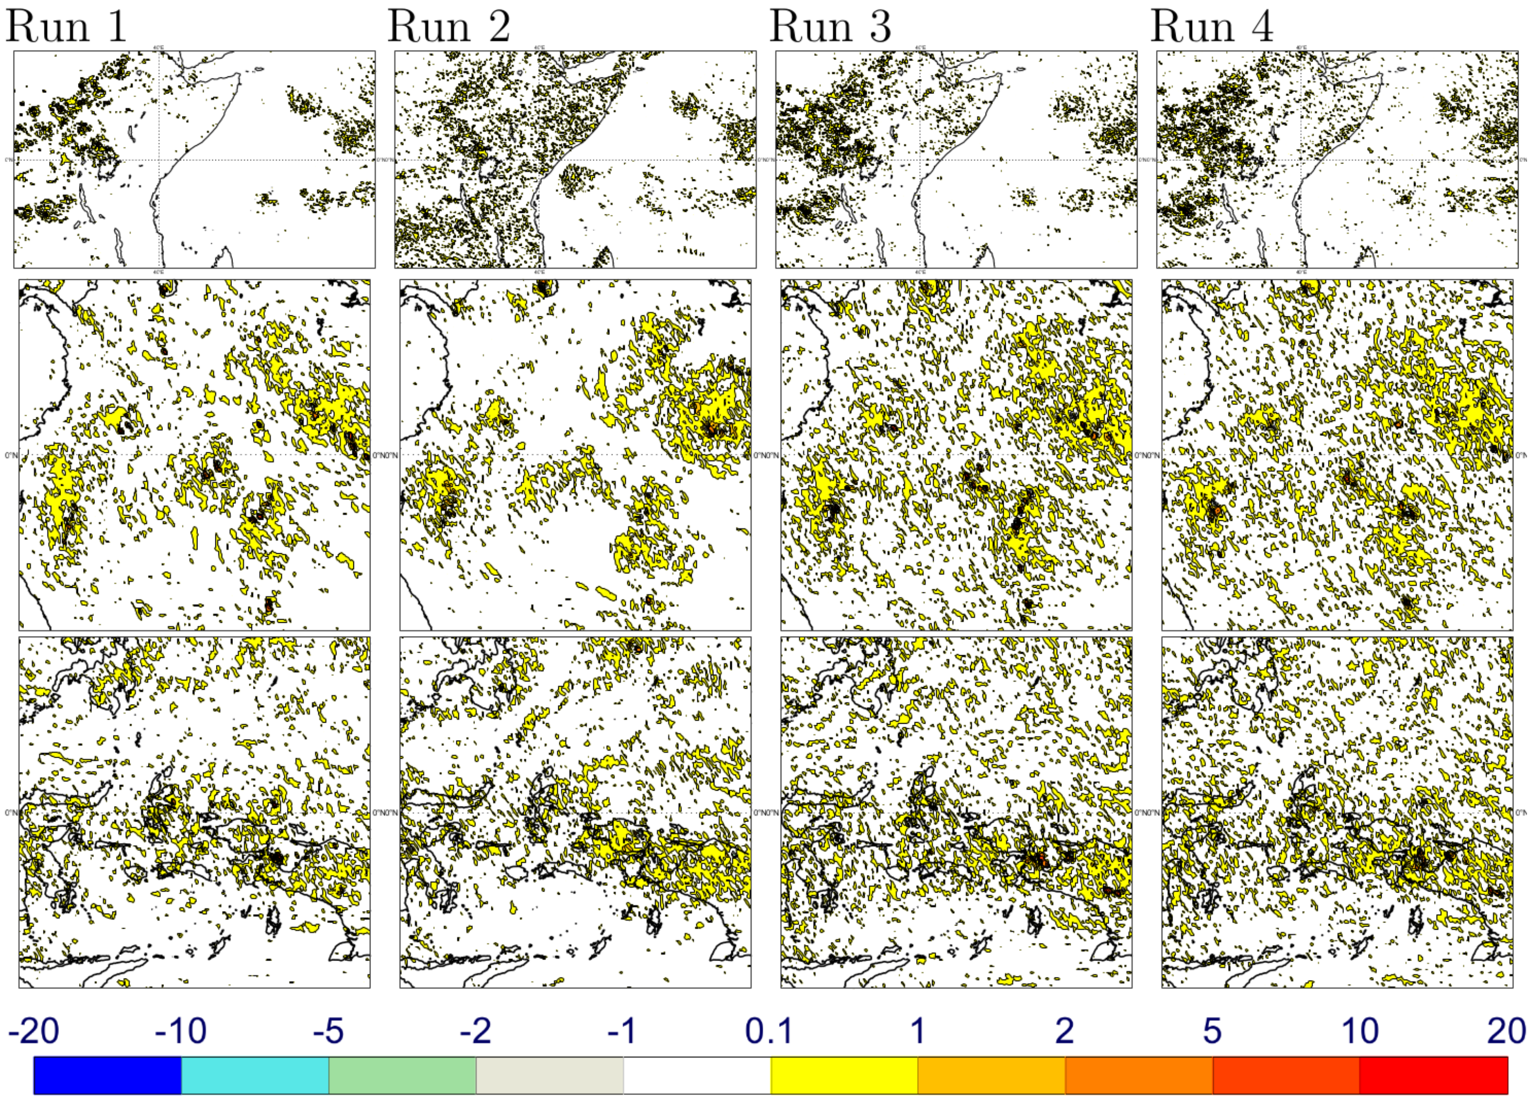

(Dueben, Wedi, Saarinen, Zeman, JMSJ 2020)

Model effects on simulating convective processes in COSMO and IFS

- *Christian Zeman et al*, *(PhD at ETH with Christoph Schär),* in preparation, e.g. Hydrostatic Approximation, Horizontal Resolution, Timestep Size, ...

PDFs of
vertical velocity
as a function of timestep

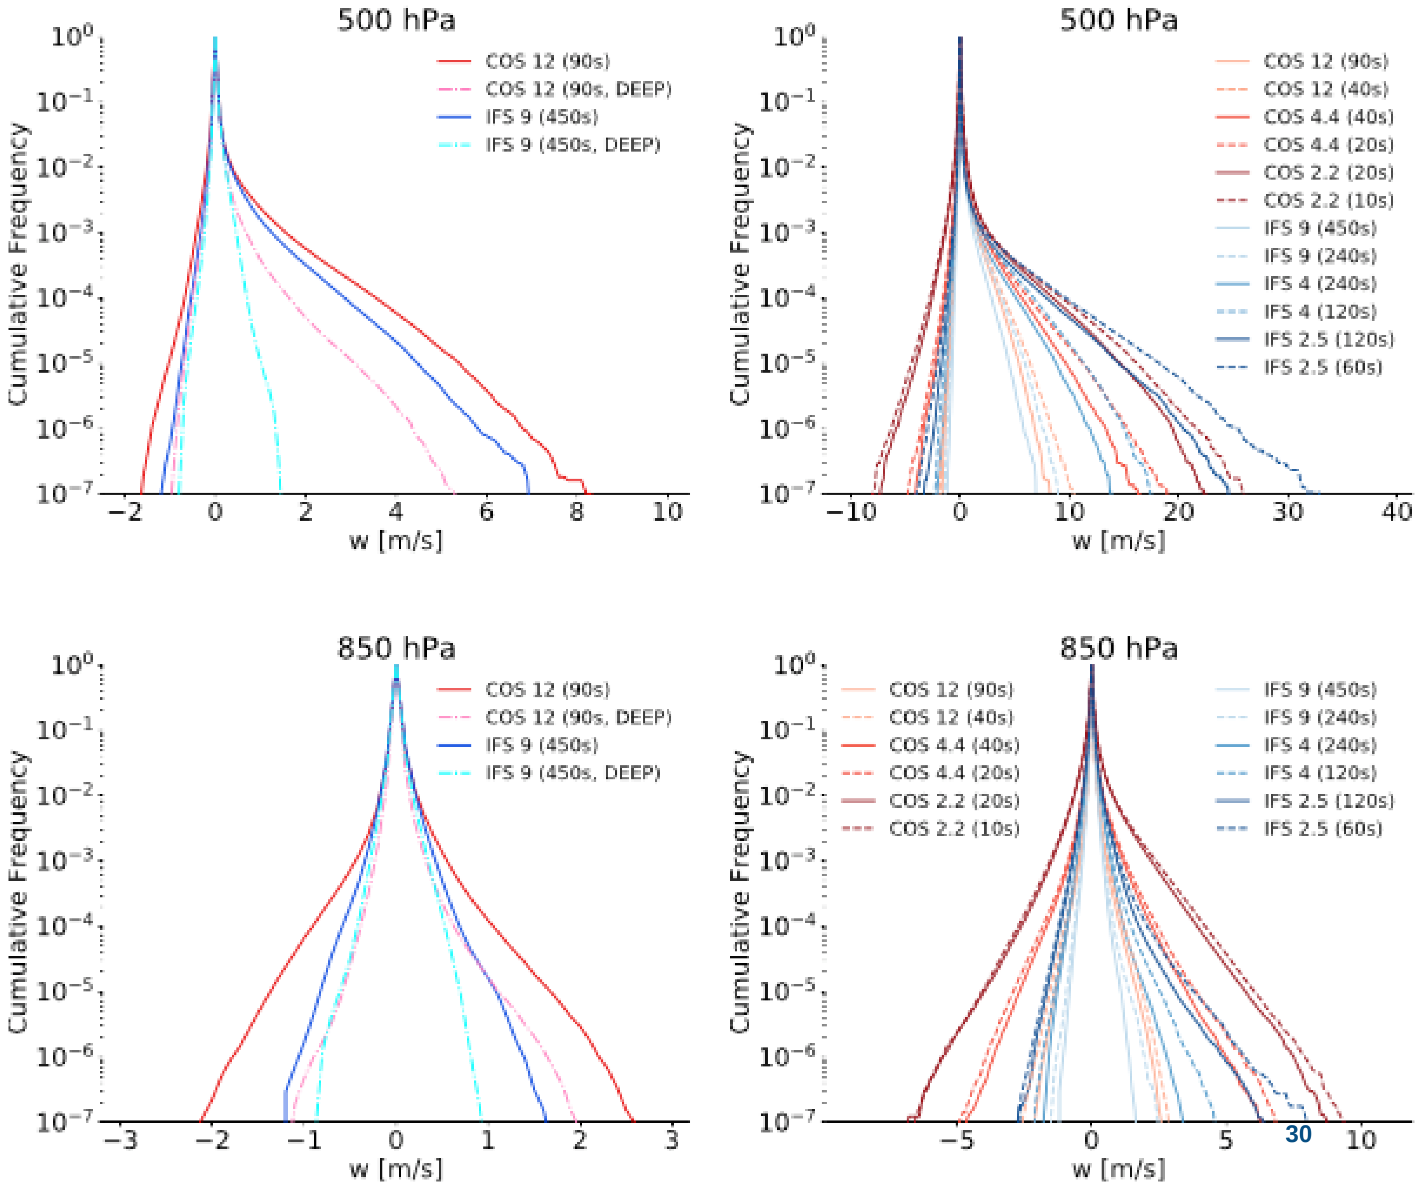

Model effects on simulating convective processes in COSMO and IFS

- *Christian Zeman et al, (PhD at ETH with Christoph Schär),* in preparation, e.g. Hydrostatic Approximation, Horizontal Resolution, Timestep Size, ...

Maps of vertical velocity

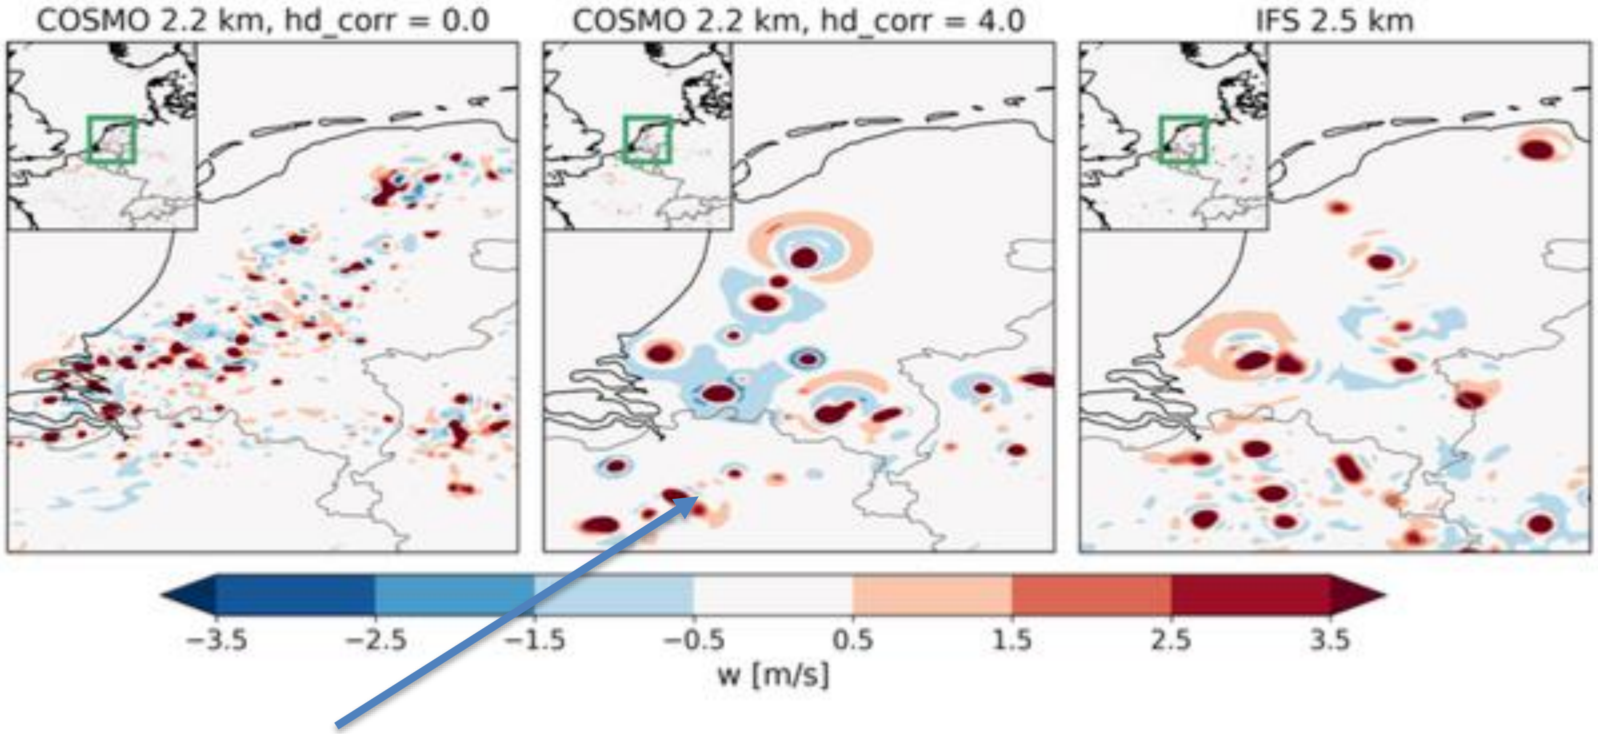

Testing numerical effects: increased horizontal diffusion == COSMO 2.2 km more similar to IFS at 2.5 km

Strong scaling at 1.4 km

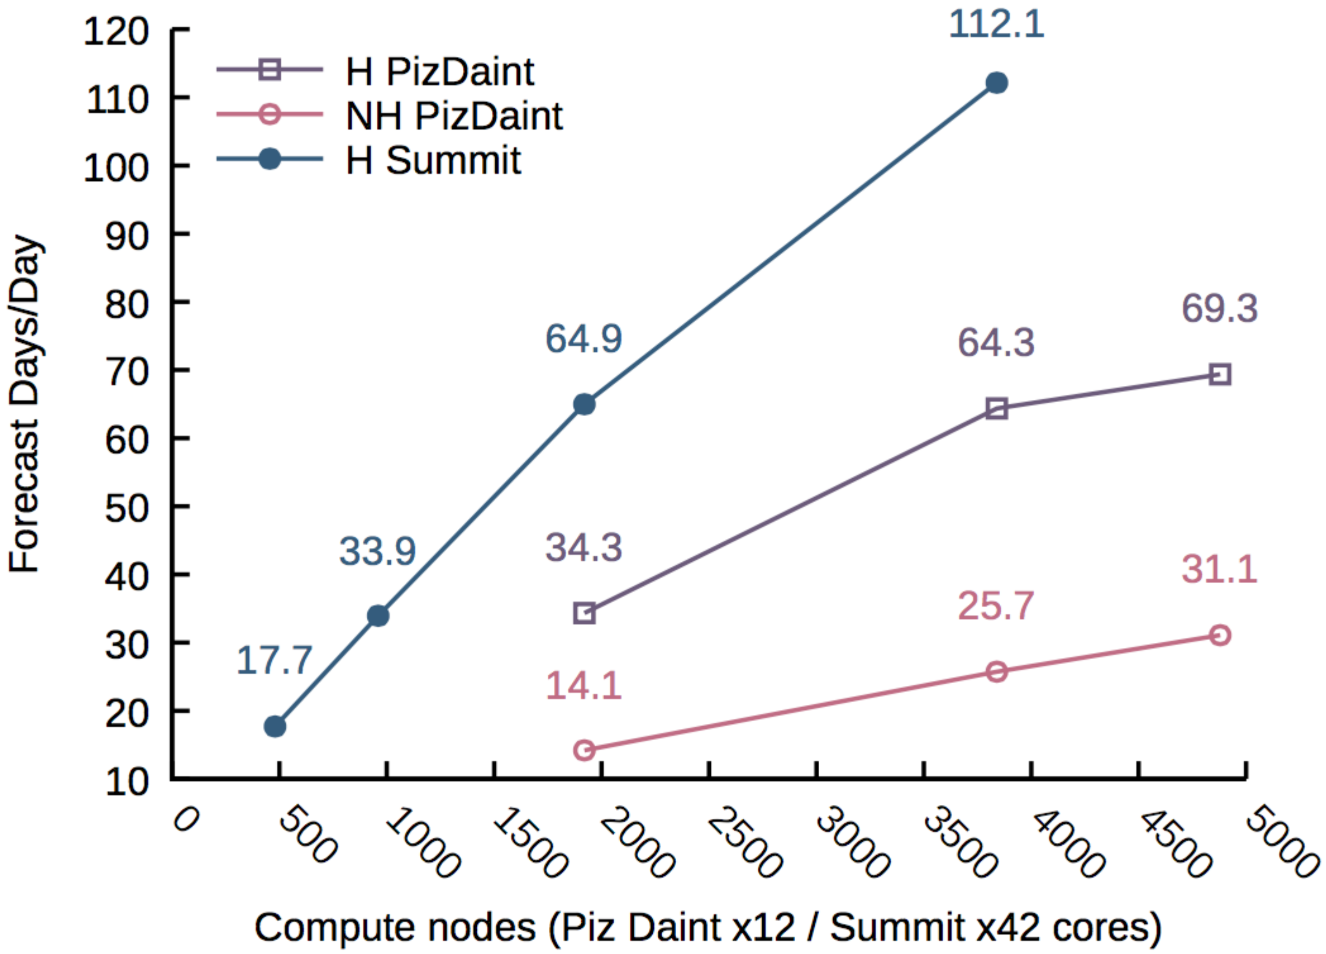

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

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)

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. Wedi[1], Inna Polichtchouk[1], Peter Dueben[1], Valentine G. Anantharaj[2], Peter Bauer[1], Souhail Boussetta[1], Philip Browne[1], Willem Deconinck[1], Wayne Gaudin[3], Ioan Hadade[1], Sam Hatfield[1], Olivier Iffrig[1], Philippe Lopez[1], Pedro Maciel[1], Andreas Mueller[1], Sami Saarinen[1], Irina Sandu[1], Tiago Quintino[1], Frederic Vitart[1] *(JAMES, 2020, accepted)*

Zonal-mean zonal wind

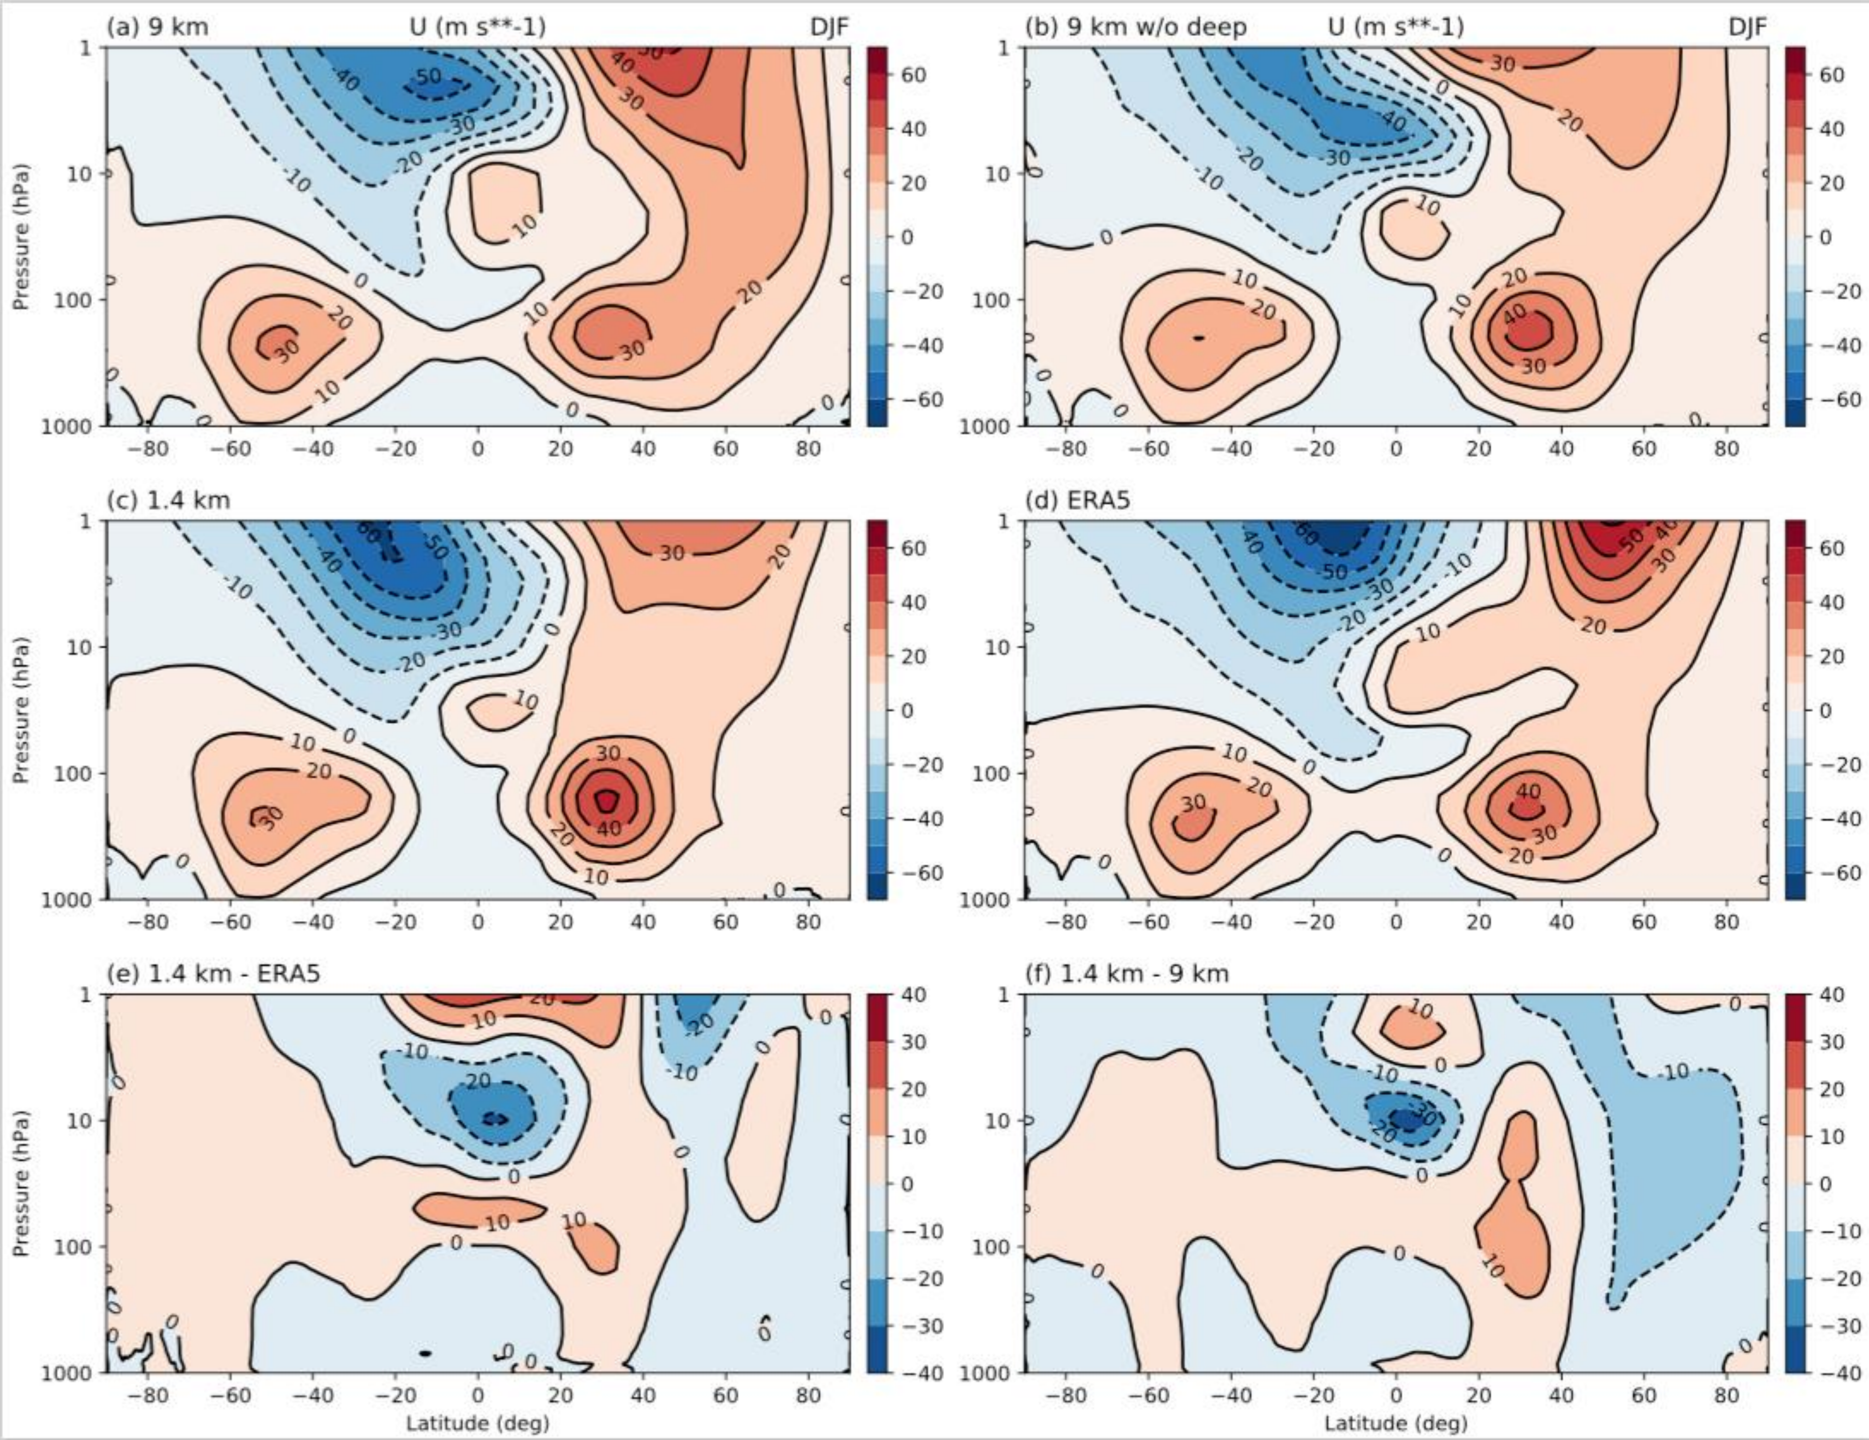

1.4 km

2018110100+48h

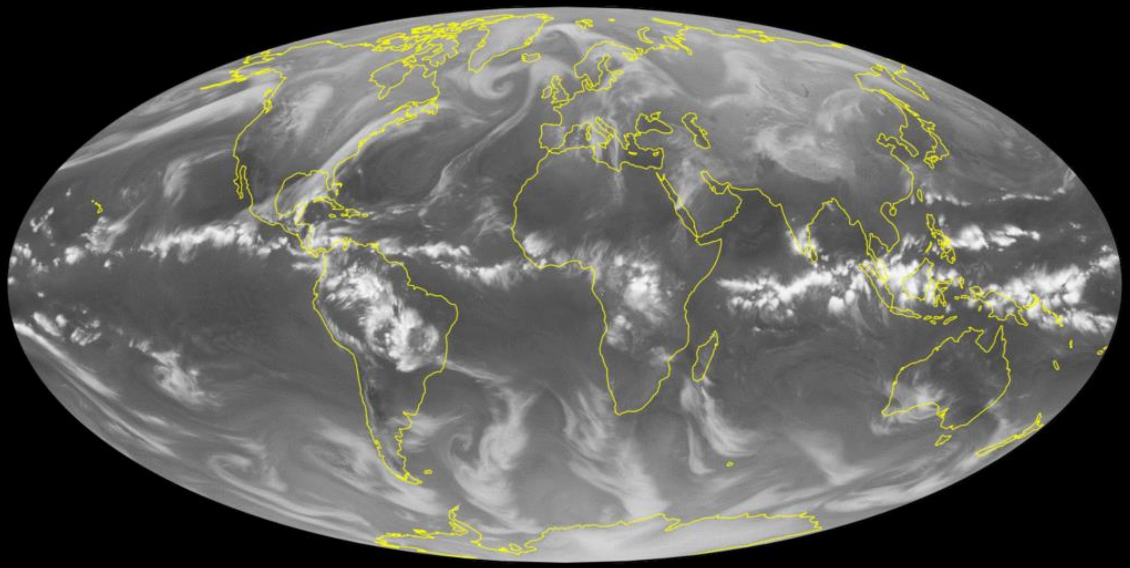

Meteosat-8, Meteosat-11 and GOES-15
2018110300

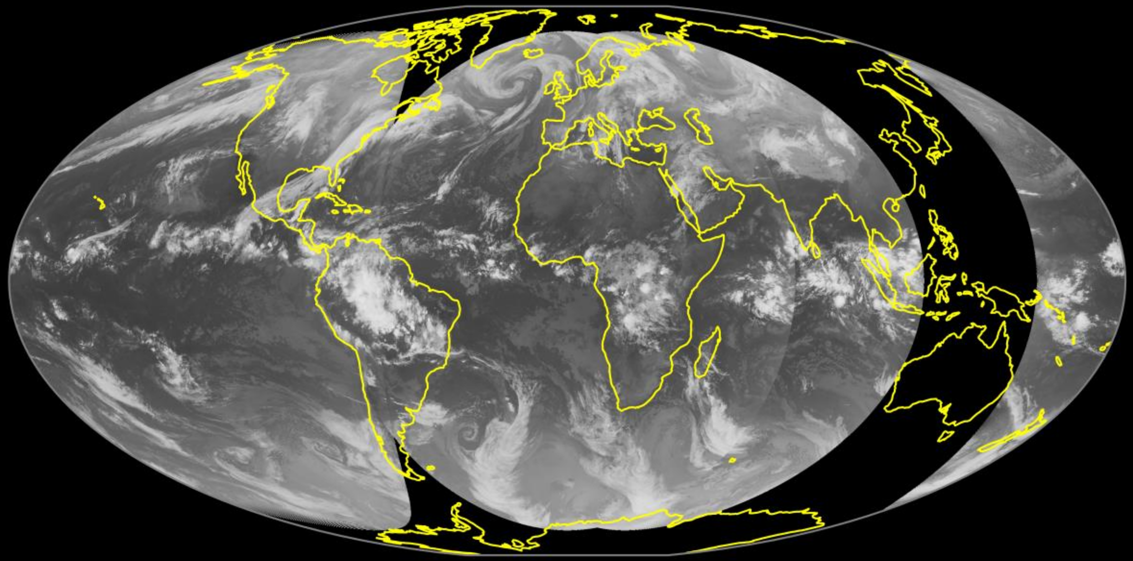

Philippe Lopez & Cristina Lupu

9 km

2018110100+48h

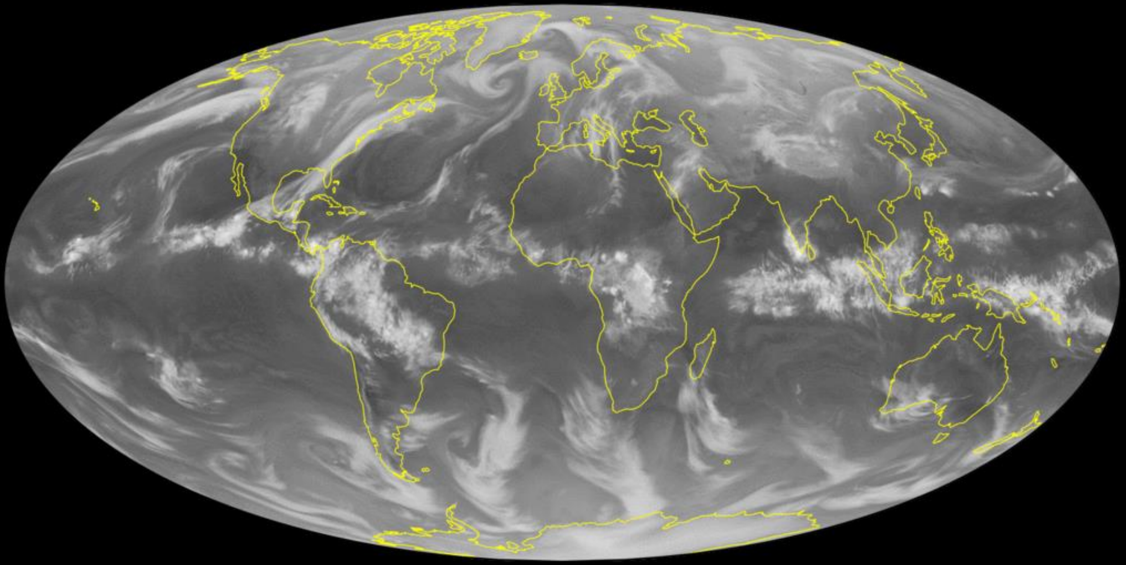

Meteosat-8, Meteosat-11 and GOES-15
2018110300

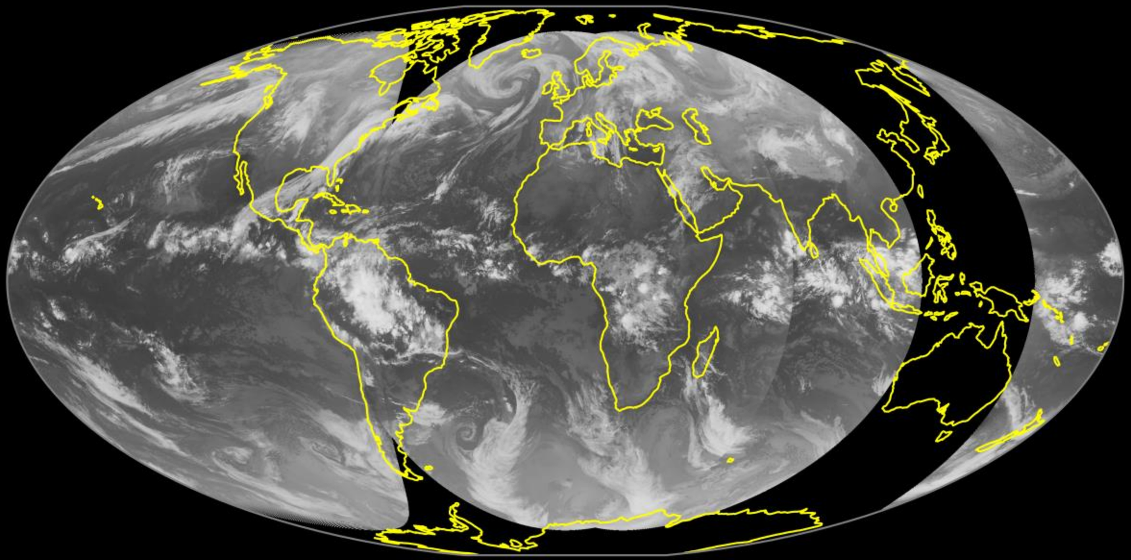

Philippe Lopez & Cristina Lupu

9 km w/o deep

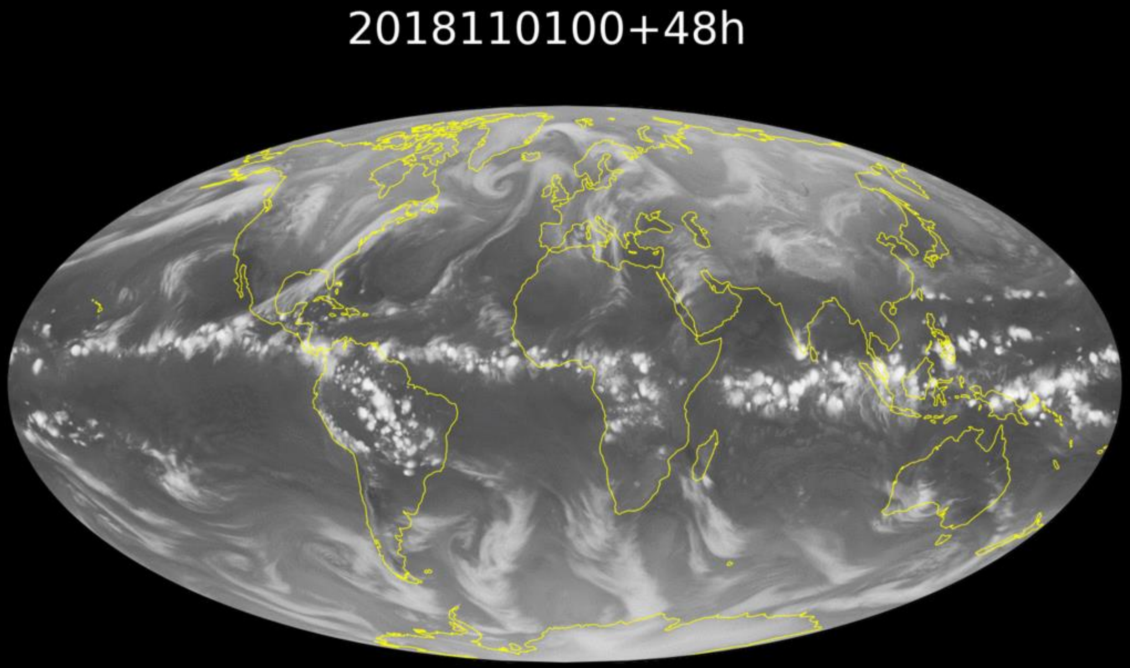

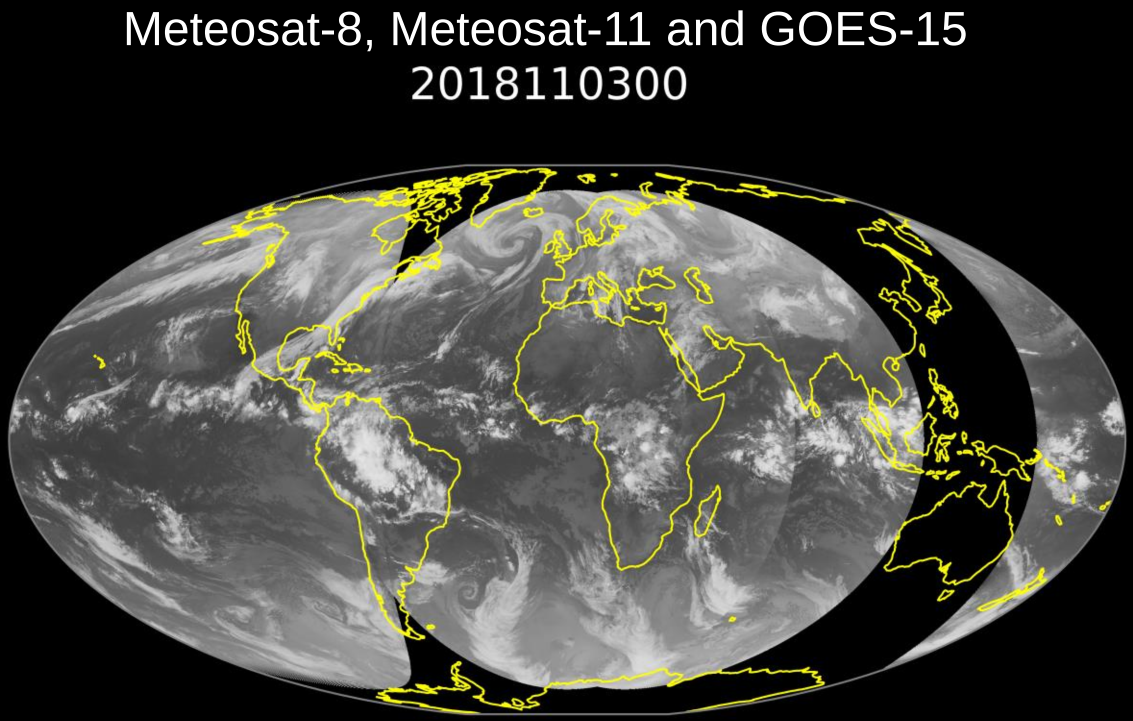

Philippe Lopez & Cristina Lupu

And what you don’t see in the satellite pics …

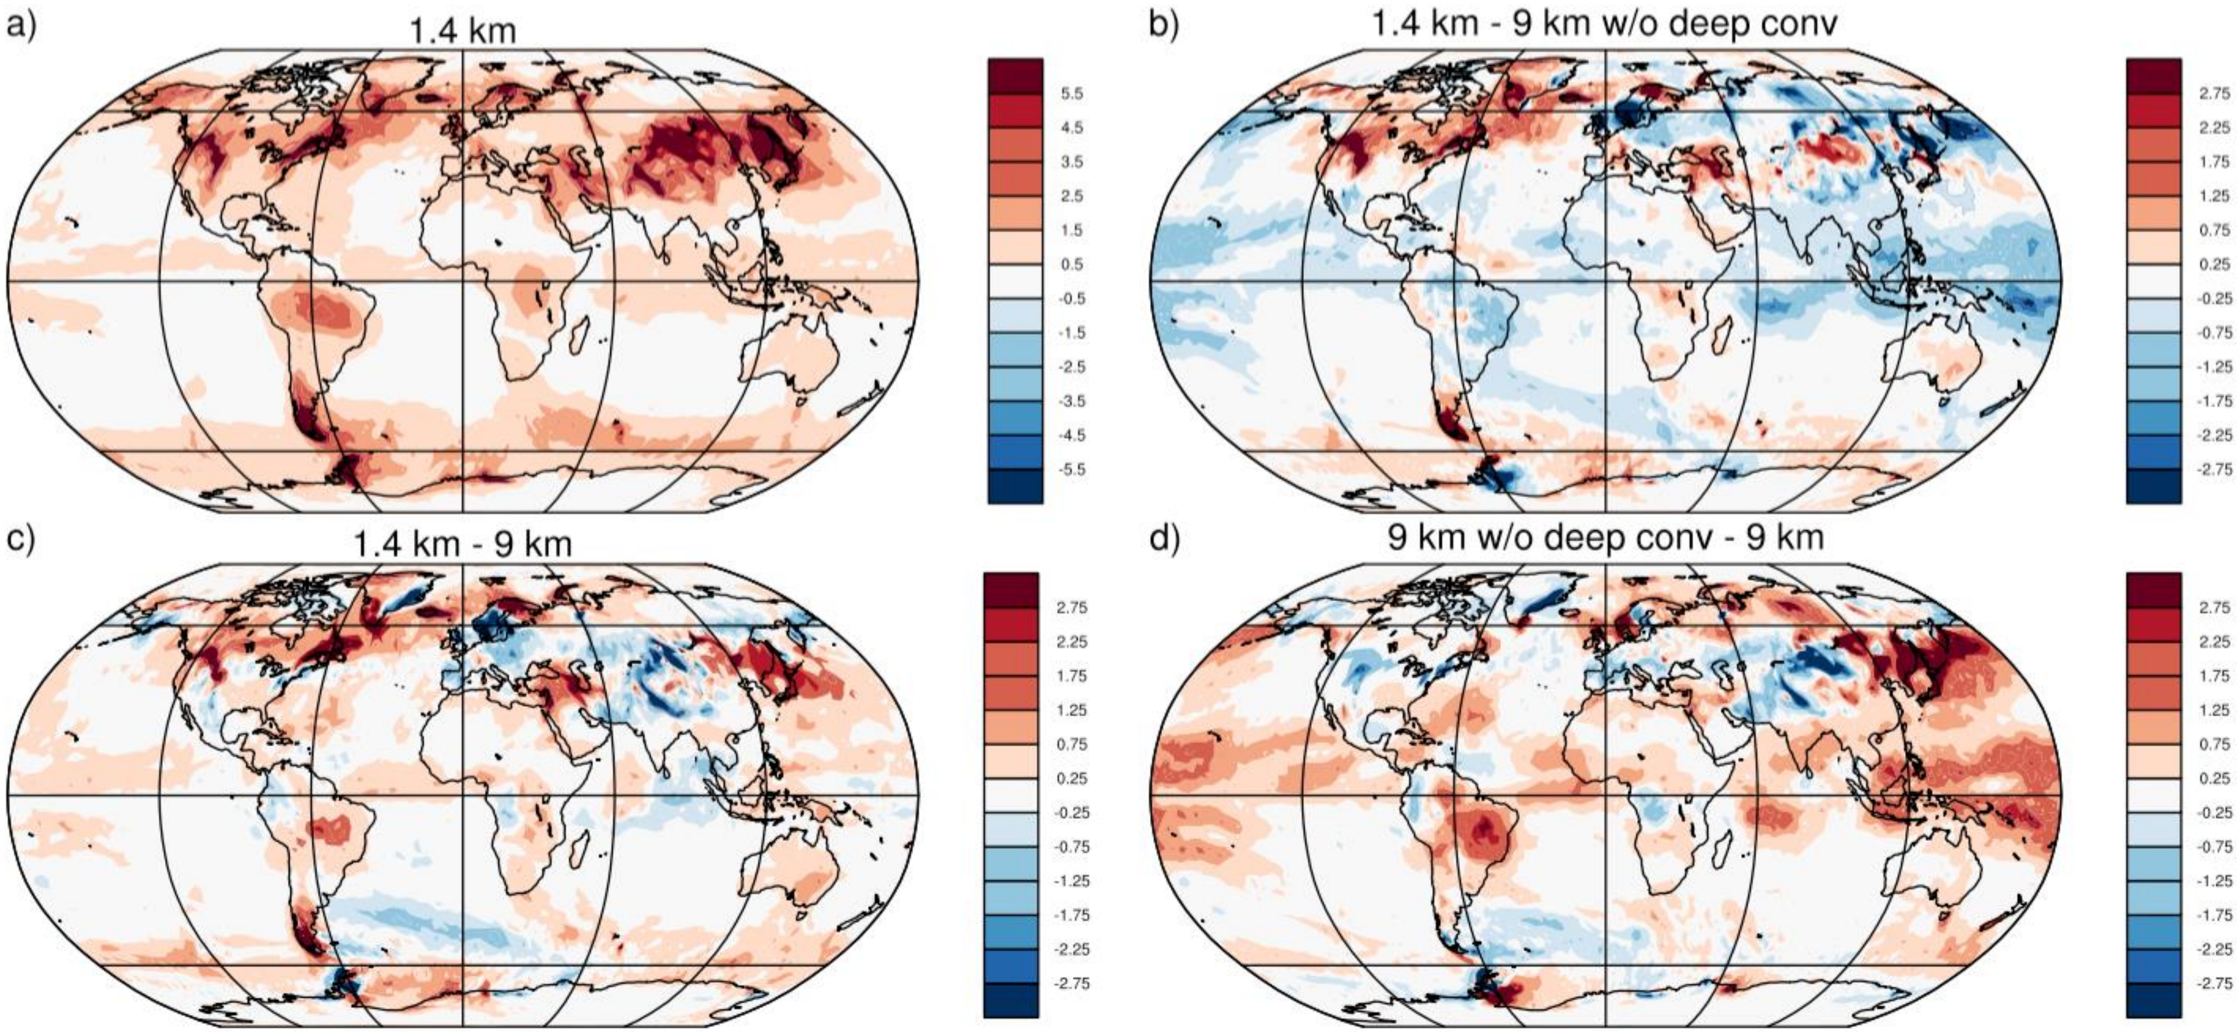

Zonal-mean absolute gravity wave momentum flux [mPa], computed from the total wave numbers 42-1279 for November 2018 at 50 hPa enhanced in convectively active regions, visibly much stronger with explicitly simulated convection at 9 km grid spacing.

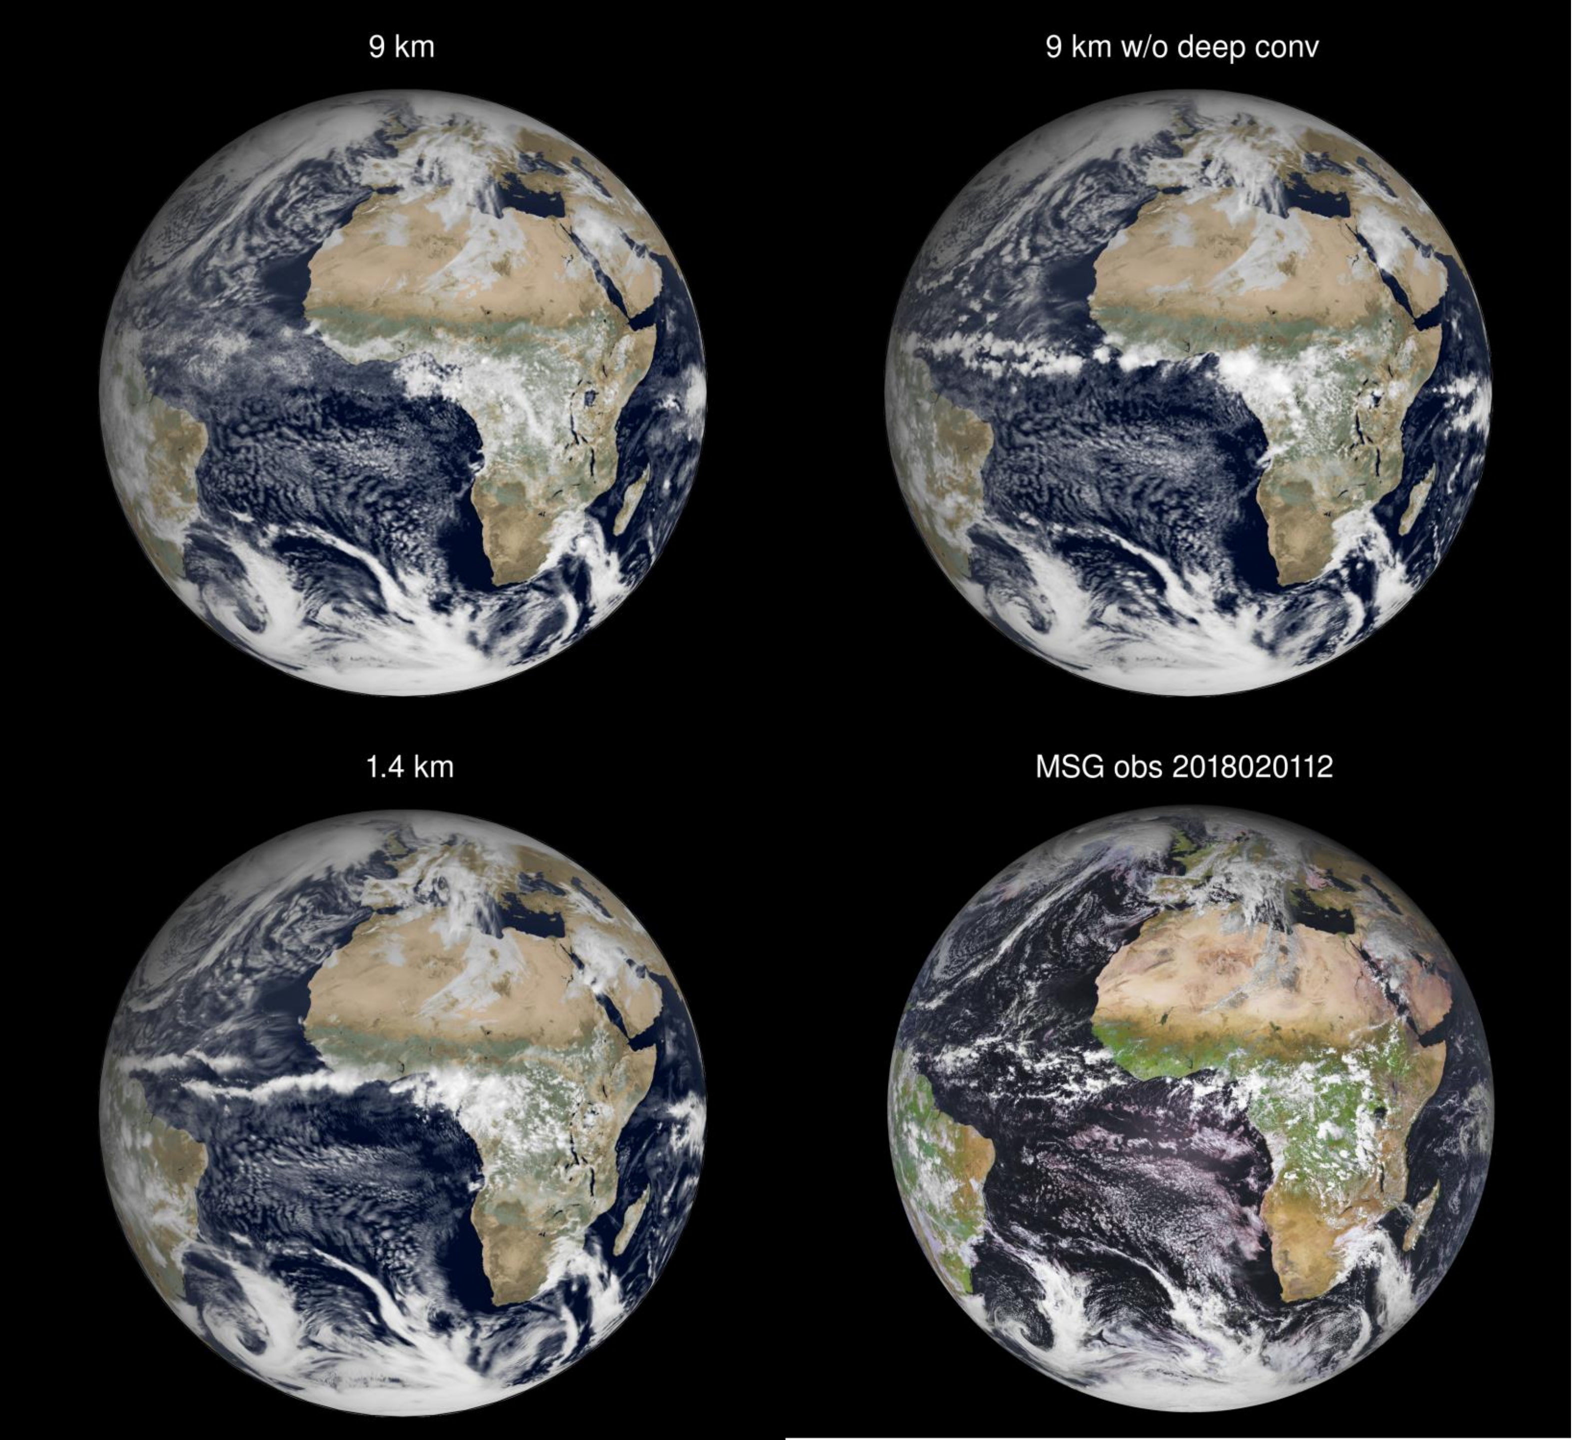

Towards a Digital Twin of the Earth System

A unified digital environment for the assessment and prediction of environmental extremes at km-scale for informed decision making at city, catchment, coastline, country and continental scale.

Digital Twin infrastructure

P. Bauer

= fairly similar observation / modelling requirements in terms extreme-scale computing and data handling!

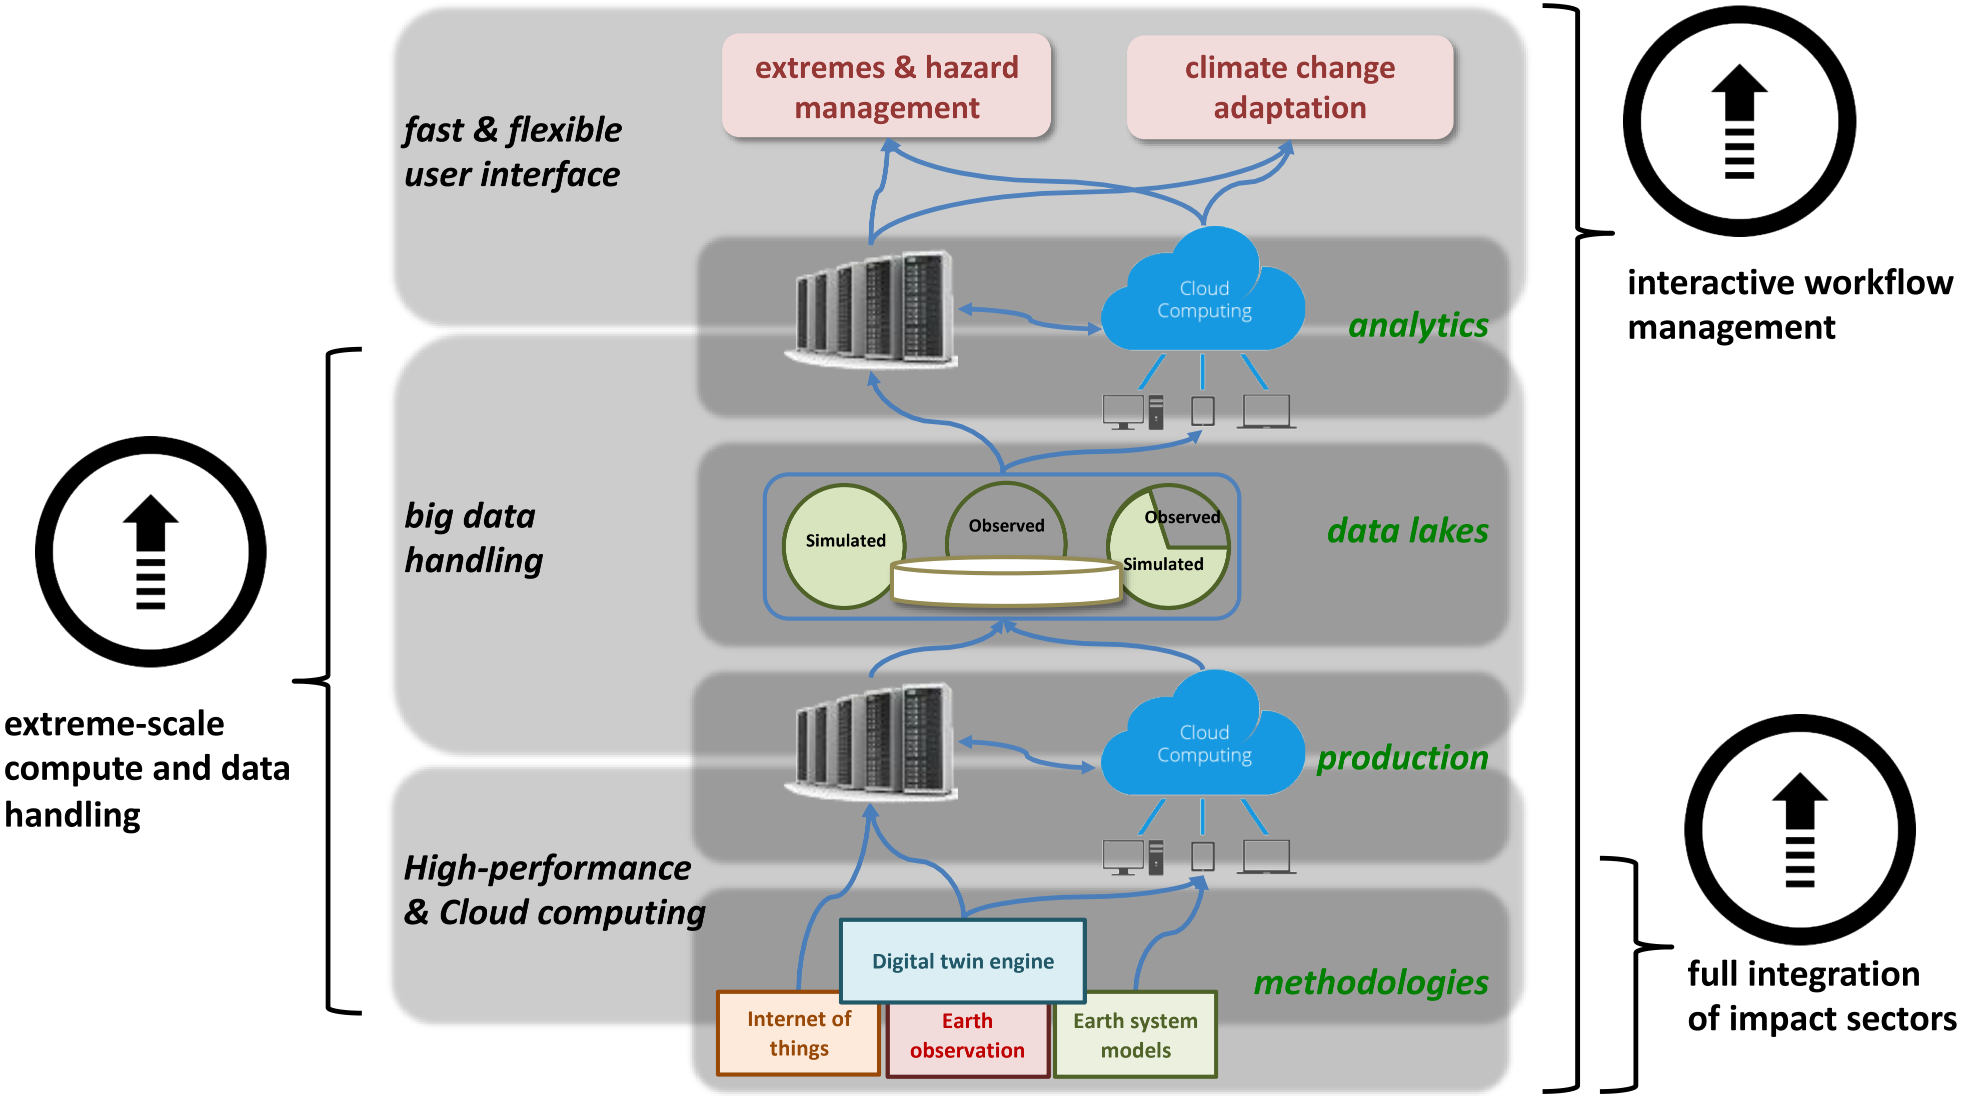

https://ec.europa.eu/digital-single-market/en/destination-earth-destine

Machine learning at ECMWF

P. Dueben

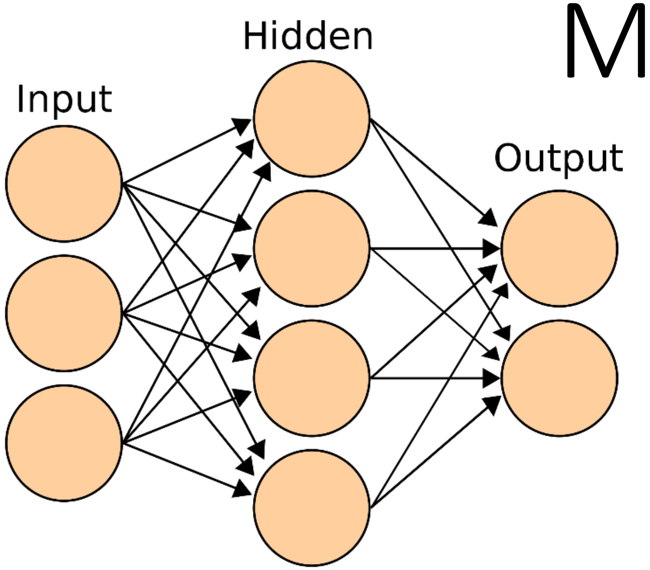

Figure copied from www.wikipedia.org

- 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
- …

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

Global 1.4 km spectra: Mid-Troposphere 500hPa

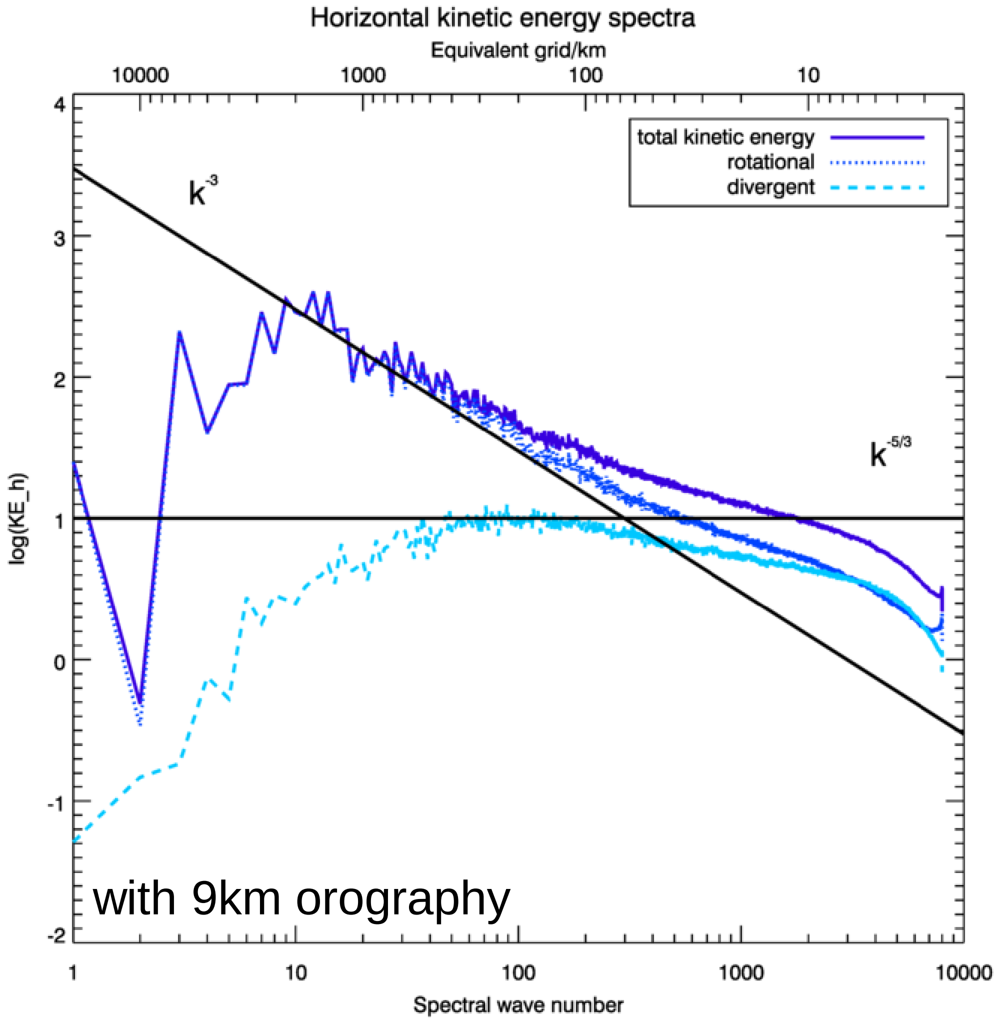

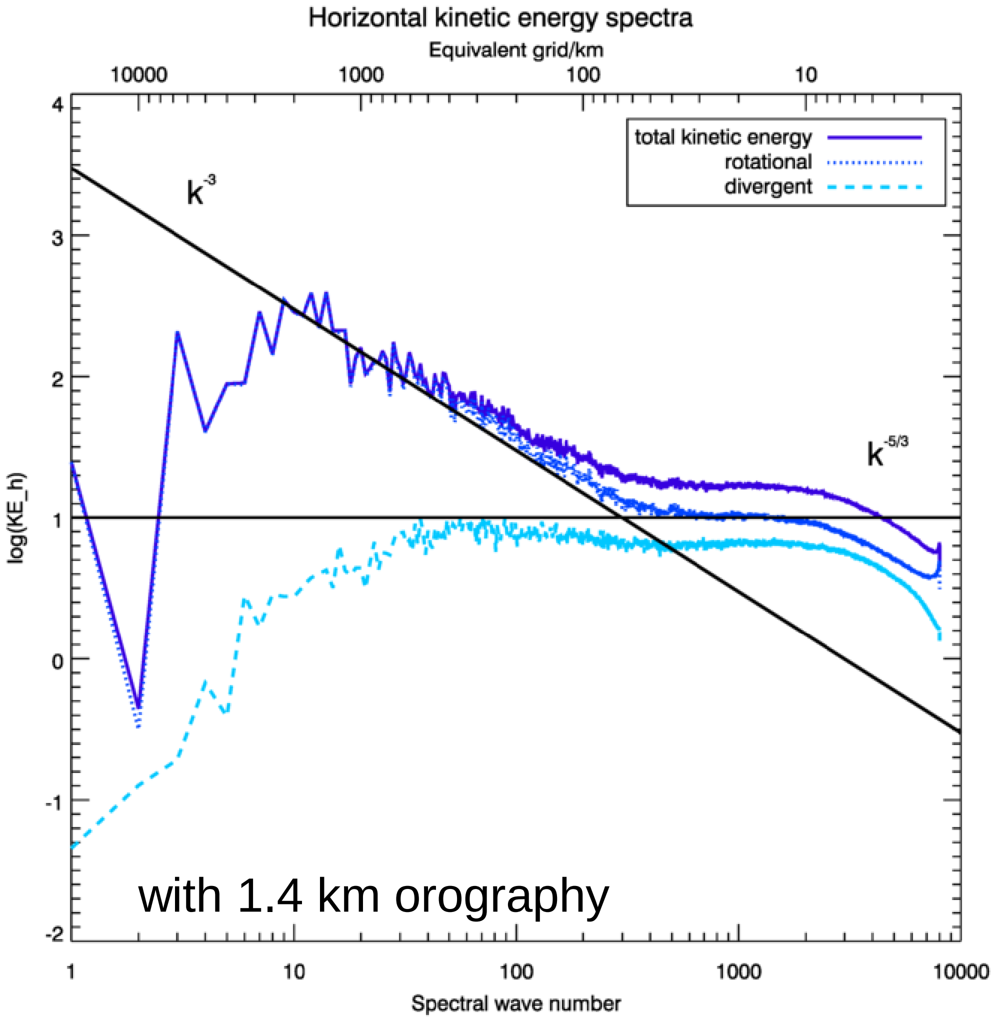

Impact of parametrizations + orography, *(Malardel + Wedi, JGR 2016)*

A vision of data-centric workflows at ECMWF

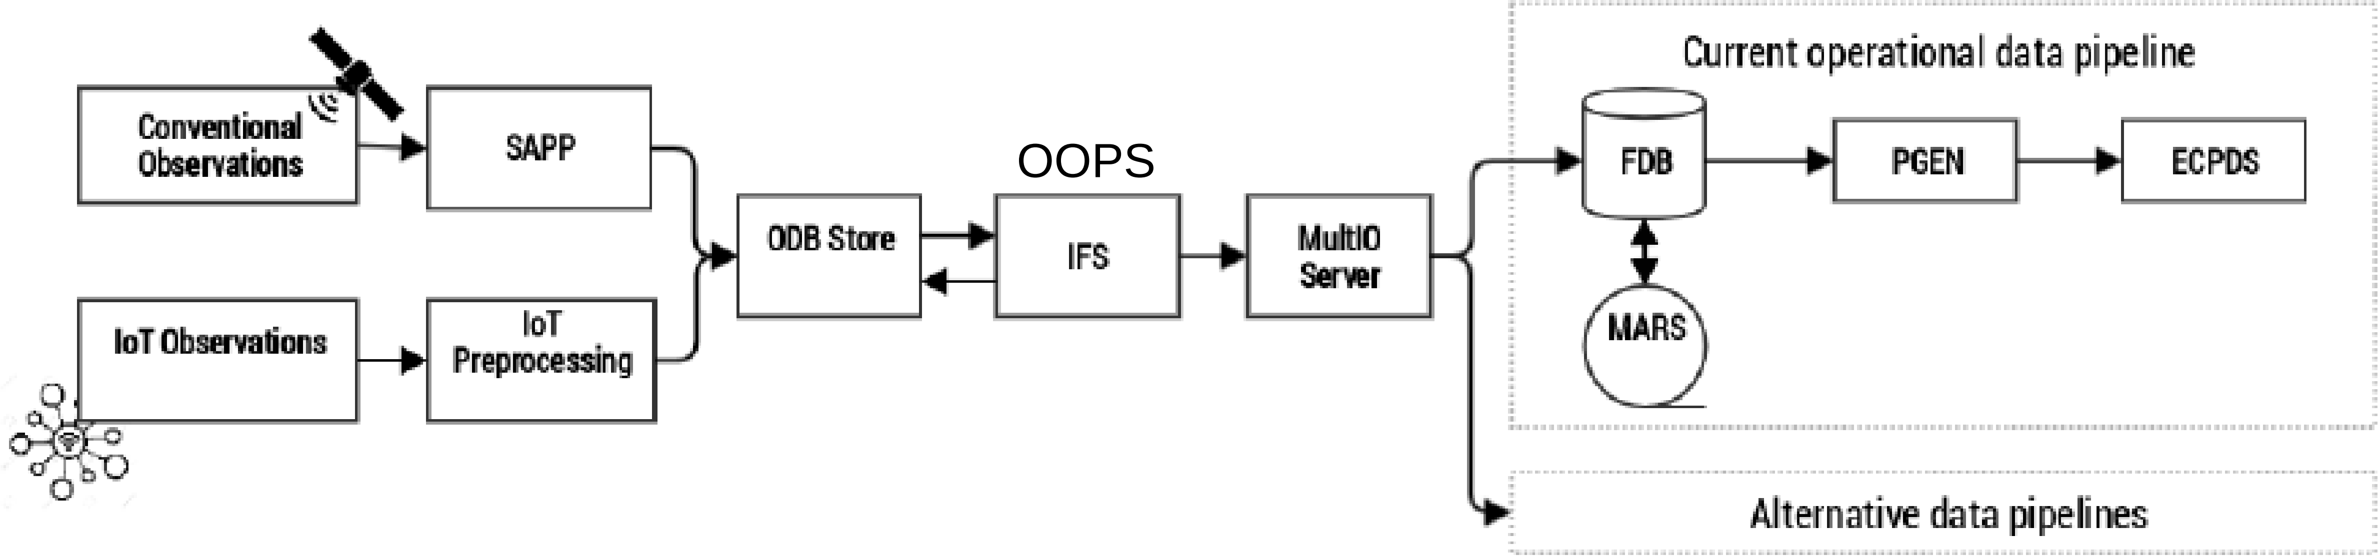

(Bauer, Quintino, et al, 2020)

- Harmonised data model beyond individual applications
- Establishing a cloud-based data handling philosophy
- Efficient user-driven data reduction
- “Polytope”, data access connecting clouds and HPC data centres