The evolving state-of-the-art in global numerical weather prediction

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

Thanks to many colleagues at ECMWF and in particular …

Further reading:

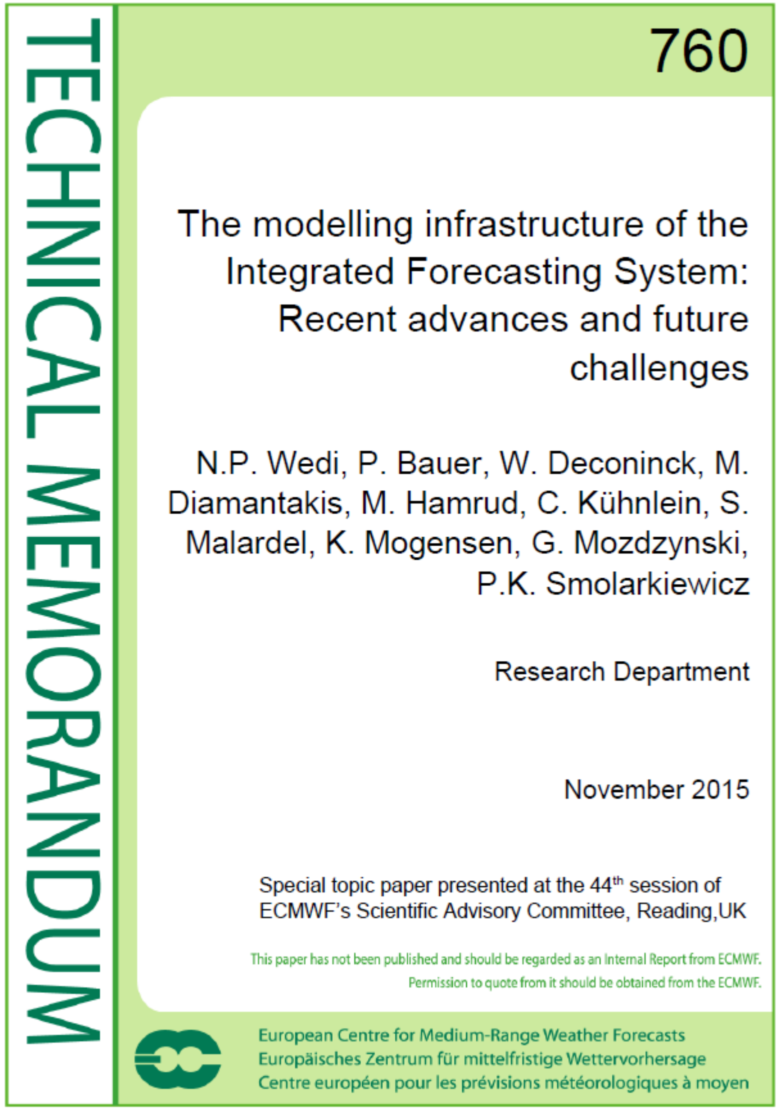

TECHNICAL MEMORANDUM

760

The modelling infrastructure of the Integrated Forecasting System: Recent advances and future challenges

N.P. Wedi, P. Bauer, W. Deconinck, M. Diamantakis, M. Hamrud, C. Kühnlein, S. Malardel, K. Mogensen, G. Mozdzynski, P.K. Smolarkiewicz

Research Department

November 2015

Special topic paper presented at the 44th session of ECMWF's Scientific Advisory Committee, Reading,UK

This paper has not been published and should be regarded as an Internal Report from ECMWF. Permission to quote from it should be obtained from the ECMWF.

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen

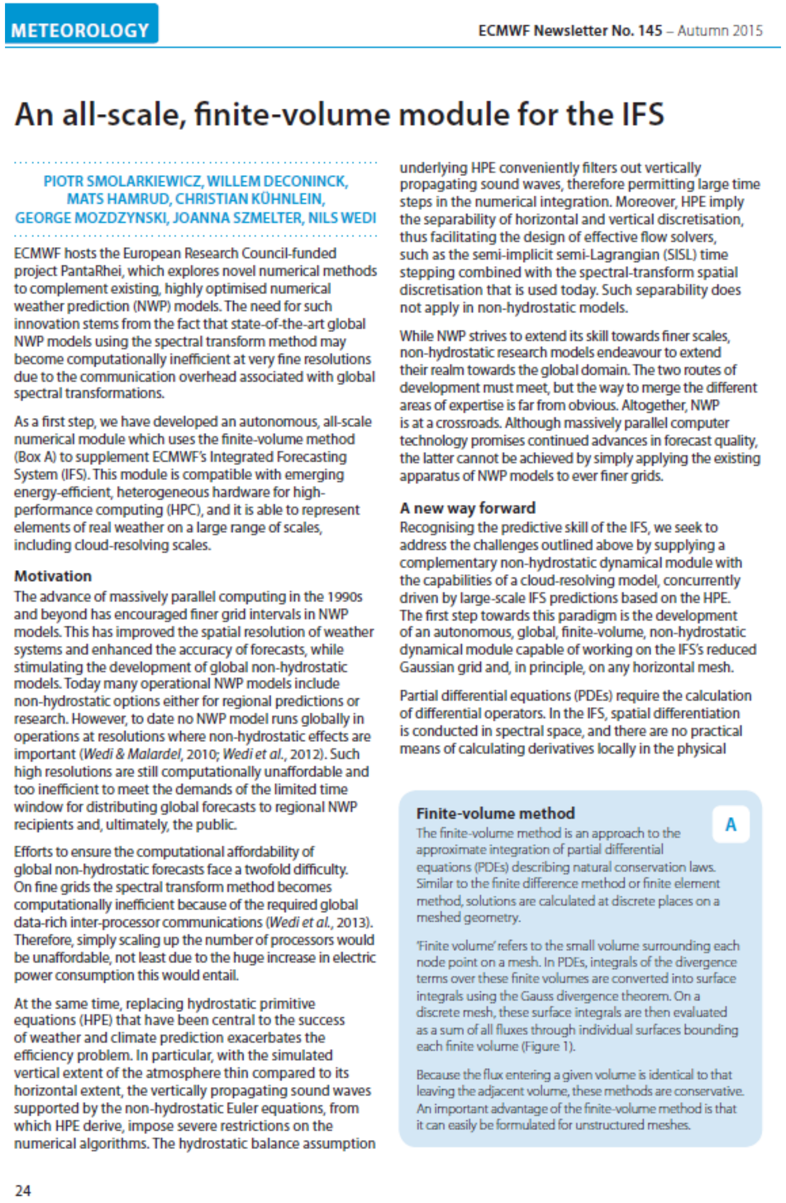

METEOROLOGY

ECMWF Newsletter No. 145 – Autumn 2015

An all-scale, finite-volume module for the IFS

PIOTR SMOLARKIEWICZ, WILLEM DECONINCK, MATS HAMRUD, CHRISTIAN KÜHNLEIN, GEORGE MOZDZYNSKI, JOANNA SZMELTER, NILS WEDI

ECMWF hosts the European Research Council-funded project PantaRhei, which explores novel numerical methods to complement existing, highly optimised numerical weather prediction (NWP) models. The need for such innovation stems from the fact that state-of-the-art global NWP models using the spectral transform method may become computationally inefficient at very fine resolutions due to the communication overhead associated with global spectral transformations.

As a first step, we have developed an autonomous, all-scale numerical module which uses the finite-volume method (Box A) to supplement ECMWF's Integrated Forecasting System (IFS). This module is compatible with emerging energy-efficient, heterogeneous hardware for high-performance computing (HPC), and it is able to represent elements of real weather on a large range of scales, including cloud-resolving scales.

Motivation

The advance of massively parallel computing in the 1990s and beyond has encouraged finer grid intervals in NWP models. This has improved the spatial resolution of weather systems and enhanced the accuracy of forecasts, while stimulating the development of global non-hydrostatic models. Today many operational NWP models include non-hydrostatic options either for regional predictions or research. However, to date no NWP model runs globally in operations at resolutions where non-hydrostatic effects are important (*Wedi & Malardel*, 2010; *Wedi et al.*, 2012). Such high resolutions are still computationally unaffordable and too inefficient to meet the demands of the limited time window for distributing global forecasts to regional NWP recipients and, ultimately, the public.

Efforts to ensure the computational affordability of global non-hydrostatic forecasts face a twofold difficulty. On fine grids the spectral transform method becomes computationally inefficient because of the required global data-rich inter-processor communications (*Wedi et al.*, 2013). Therefore, simply scaling up the number of processors would be unaffordable, not least due to the huge increase in electric power consumption this would entail.

At the same time, replacing hydrostatic primitive equations (HPE) that have been central to the success of weather and climate prediction exacerbates the efficiency problem. In particular, with the simulated vertical extent of the atmosphere thin compared to its horizontal extent, the vertically propagating sound waves supported by the non-hydrostatic Euler equations, from which HPE derive, impose severe restrictions on the numerical algorithms. The hydrostatic balance assumption underlying HPE conveniently filters out vertically propagating sound waves, therefore permitting large time steps in the numerical integration. Moreover, HPE imply the separability of horizontal and vertical discretisation, thus facilitating the design of effective flow solvers, such as the semi-implicit semi-Lagrangian (SISL) time stepping combined with the spectral-transform spatial discretisation that is used today. Such separability does not apply in non-hydrostatic models.

While NWP strives to extend its skill towards finer scales, non-hydrostatic research models endeavour to extend their realm towards the global domain. The two routes of development must meet, but the way to merge the different areas of expertise is far from obvious. Altogether, NWP is at a crossroads. Although massively parallel computer technology promises continued advances in forecast quality, the latter cannot be achieved by simply applying the existing apparatus of NWP models to ever finer grids.

A new way forward

Recognising the predictive skill of the IFS, we seek to address the challenges outlined above by supplying a complementary non-hydrostatic dynamical module with the capabilities of a cloud-resolving model, concurrently driven by large-scale IFS predictions based on the HPE. The first step towards this paradigm is the development of an autonomous, global, finite-volume, non-hydrostatic dynamical module capable of working on the IFS's reduced Gaussian grid and, in principle, on any horizontal mesh.

Partial differential equations (PDEs) require the calculation of differential operators. In the IFS, spatial differentiation is conducted in spectral space, and there are no practical means of calculating derivatives locally in the physical

Finite-volume method A

The finite-volume method is an approach to the approximate integration of partial differential equations (PDEs) describing natural conservation laws. Similar to the finite difference method or finite element method, solutions are calculated at discrete places on a meshed geometry.

'Finite volume' refers to the small volume surrounding each node point on a mesh. In PDEs, integrals of the divergence terms over these finite volumes are converted into surface integrals using the Gauss divergence theorem. On a discrete mesh, these surface integrals are then evaluated as a sum of all fluxes through individual surfaces bounding each finite volume (Figure 1).

Because the flux entering a given volume is identical to that leaving the adjacent volume, these methods are conservative. An important advantage of the finite-volume method is that it can easily be formulated for unstructured meshes.

24

Funded by the European Union

PantaRhei *http://www.ecmwf.int/*

European Research Council
Established by the European Commission
Supporting top researchers from anywhere in the world

P. Smolarkiewicz

ESCAPE *http://www.hpc-escape.eu/*

G. Mengaldo, A. Mueller, ESCAPE partners, …

ESiWACE *https://www.esiwace.eu/*

P. Dueben

Outline

- ECMWF’s integrated forecasting system (IFS)
  - *Global spectral transforms and a new grid*
  - *Communication cost*
- A flexible, scalable and sustainable model infrastructure
  - Algorithmic flexibility
- A new 256 Megapixel digital (Earth) camera called IFS …

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

No grid-staggering of prognostic variables

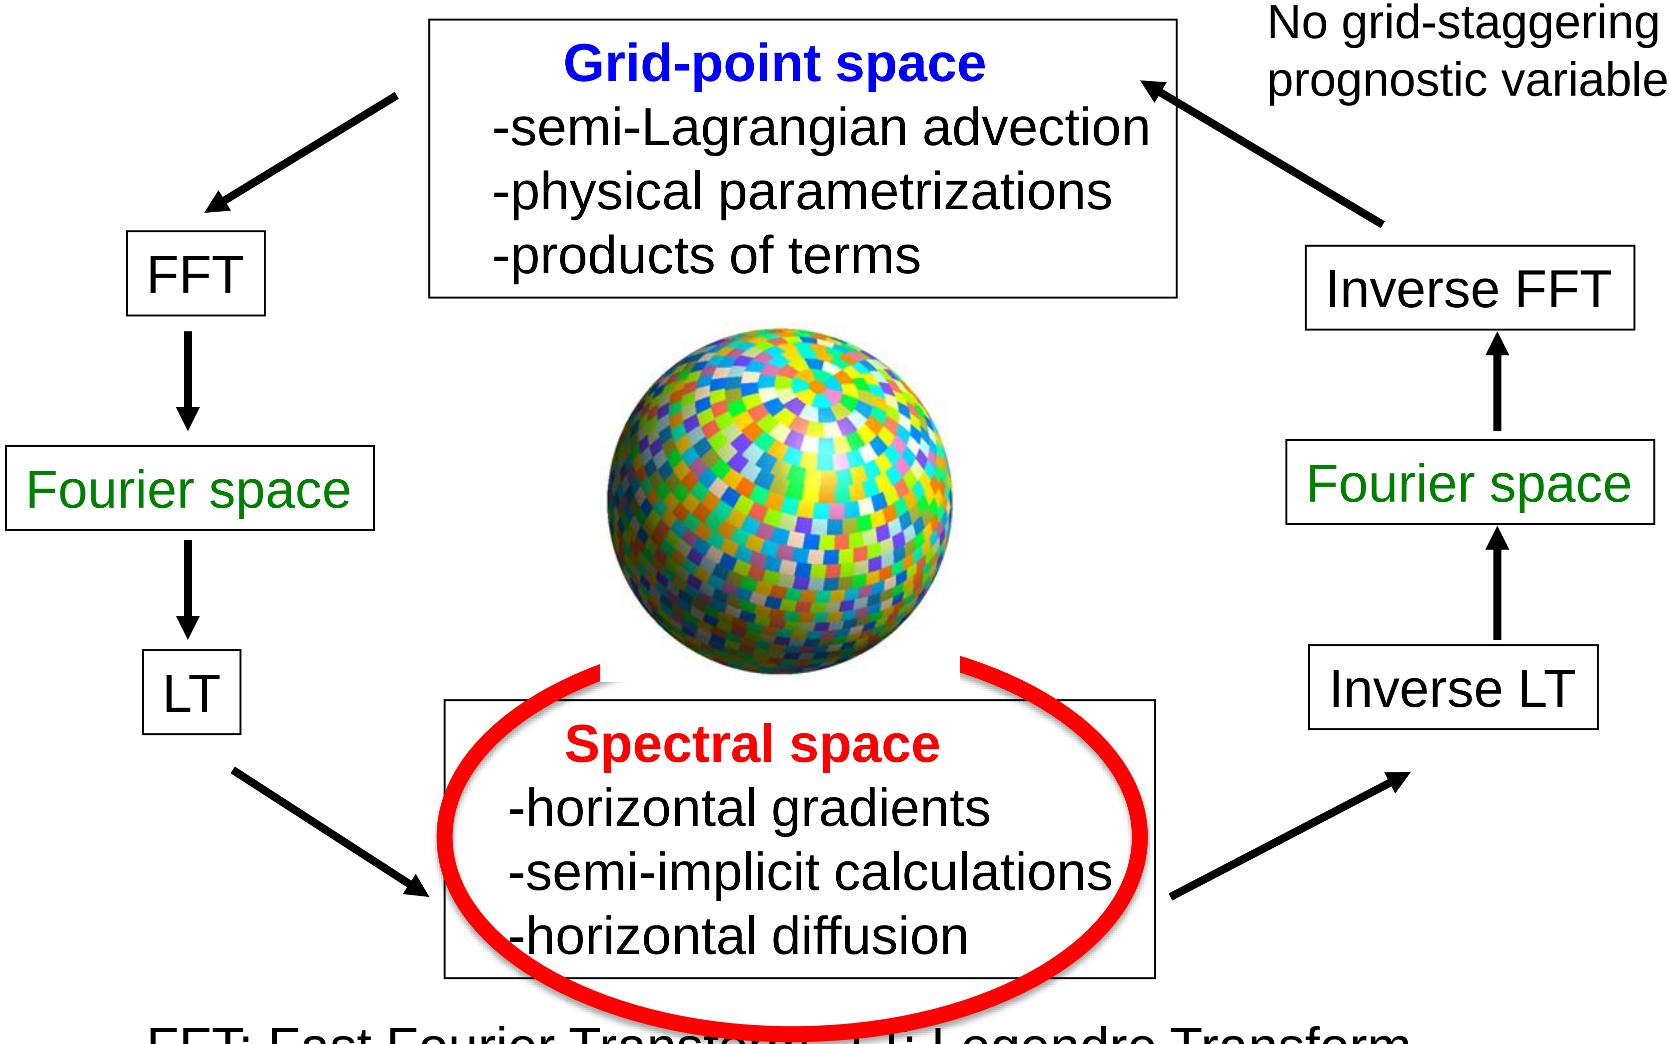

Grid-point space
- -semi-Lagrangian advection
- -physical parametrizations
- -products of terms

FFT

Fourier space

LT

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

Inverse LT

Fourier space

Inverse FFT

FFT: Fast Fourier Transform, LT: Legendre Transform

A new grid for ECMWF

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

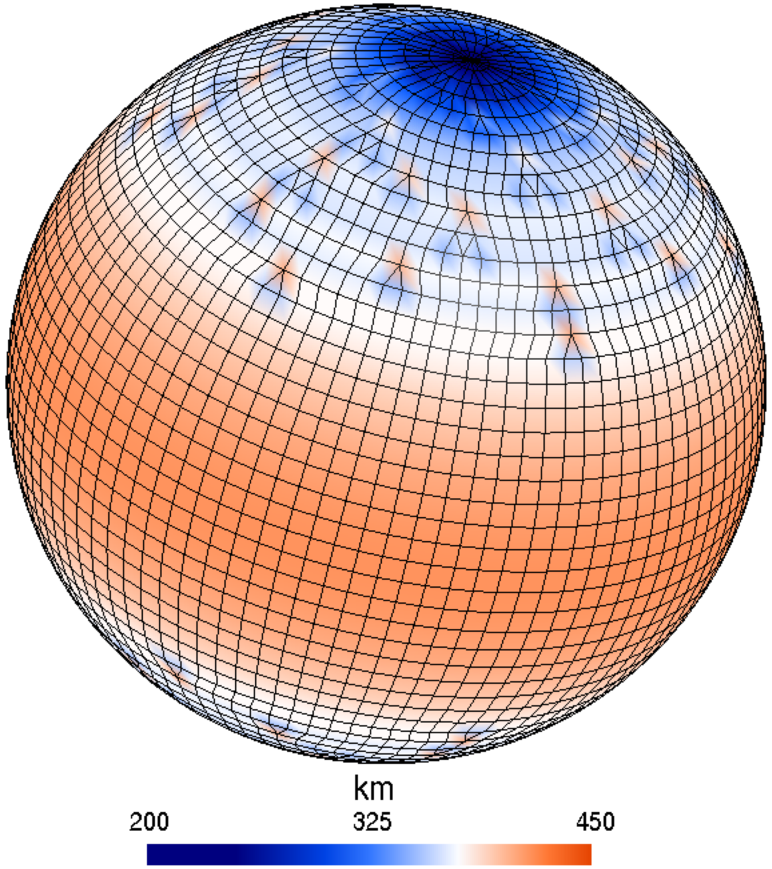

N24 reduced Gaussian grid

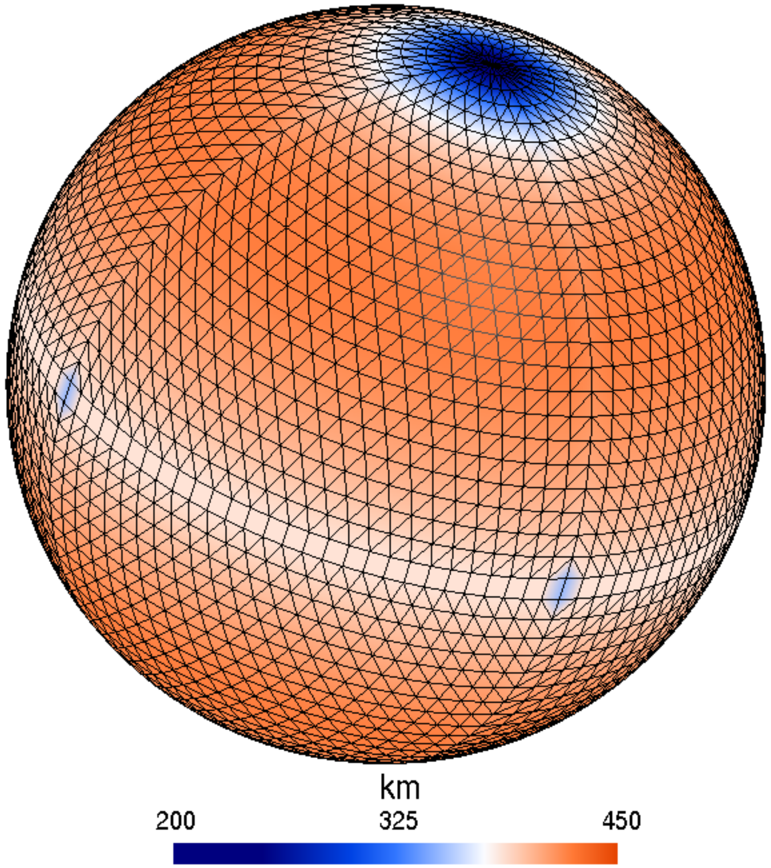

N24 octahedral Gaussian grid

(Wedi et al, 2015)

RAPS14 performance

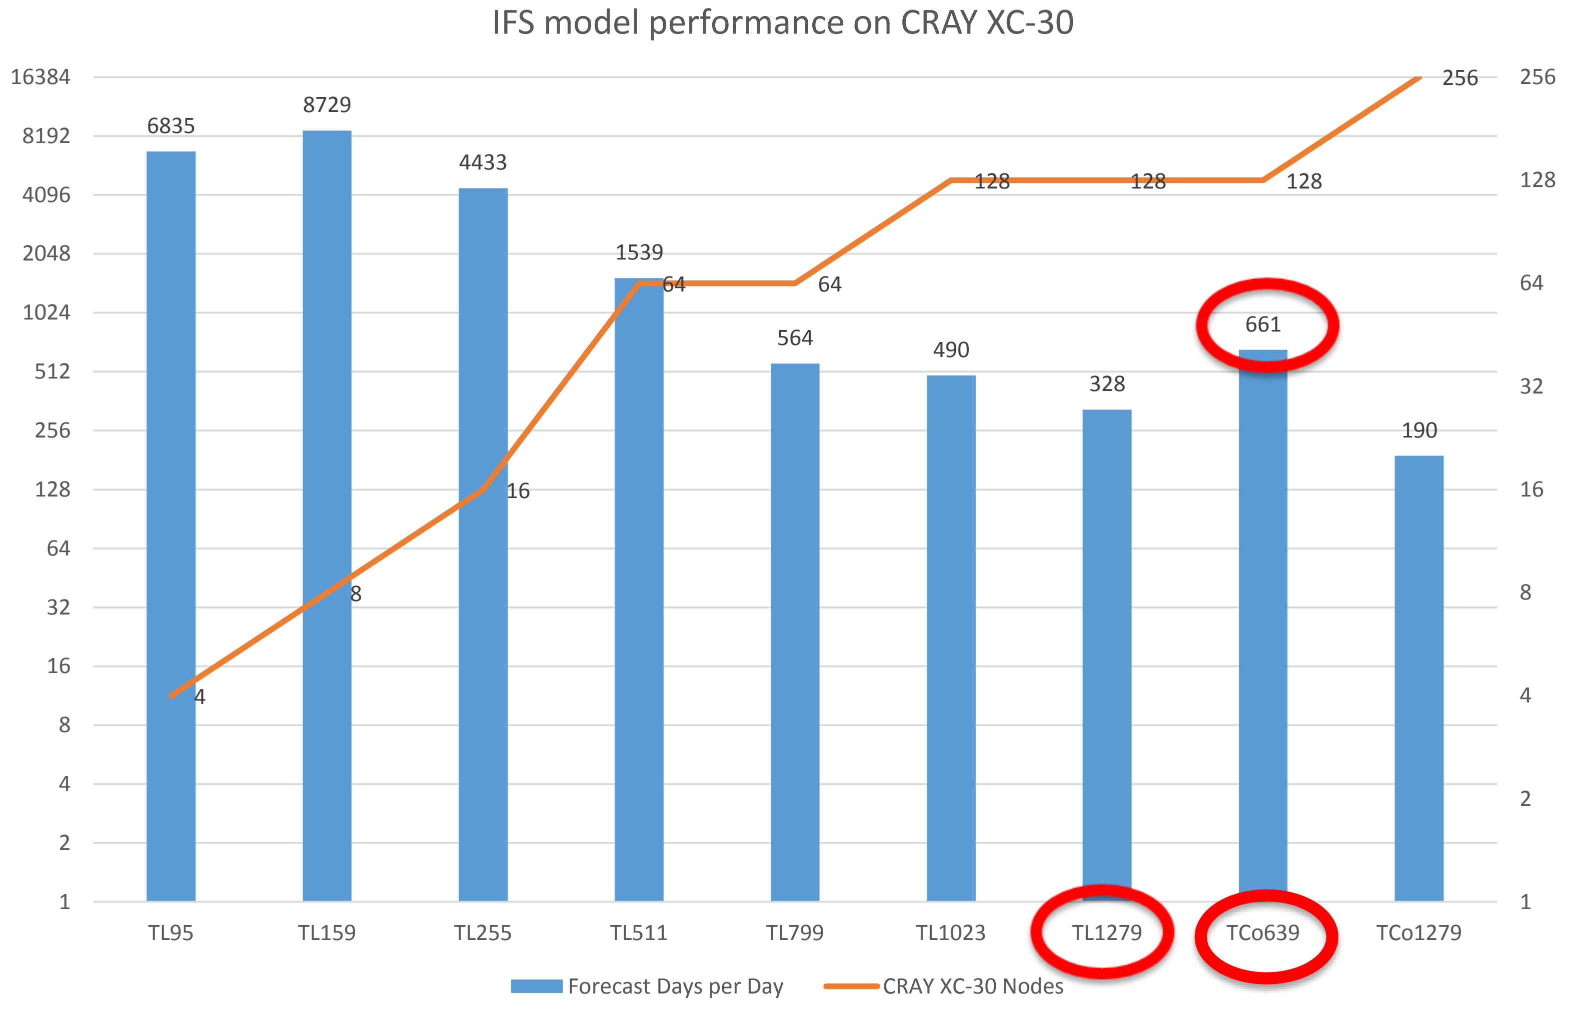

EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

MPI communication cost at large core counts …

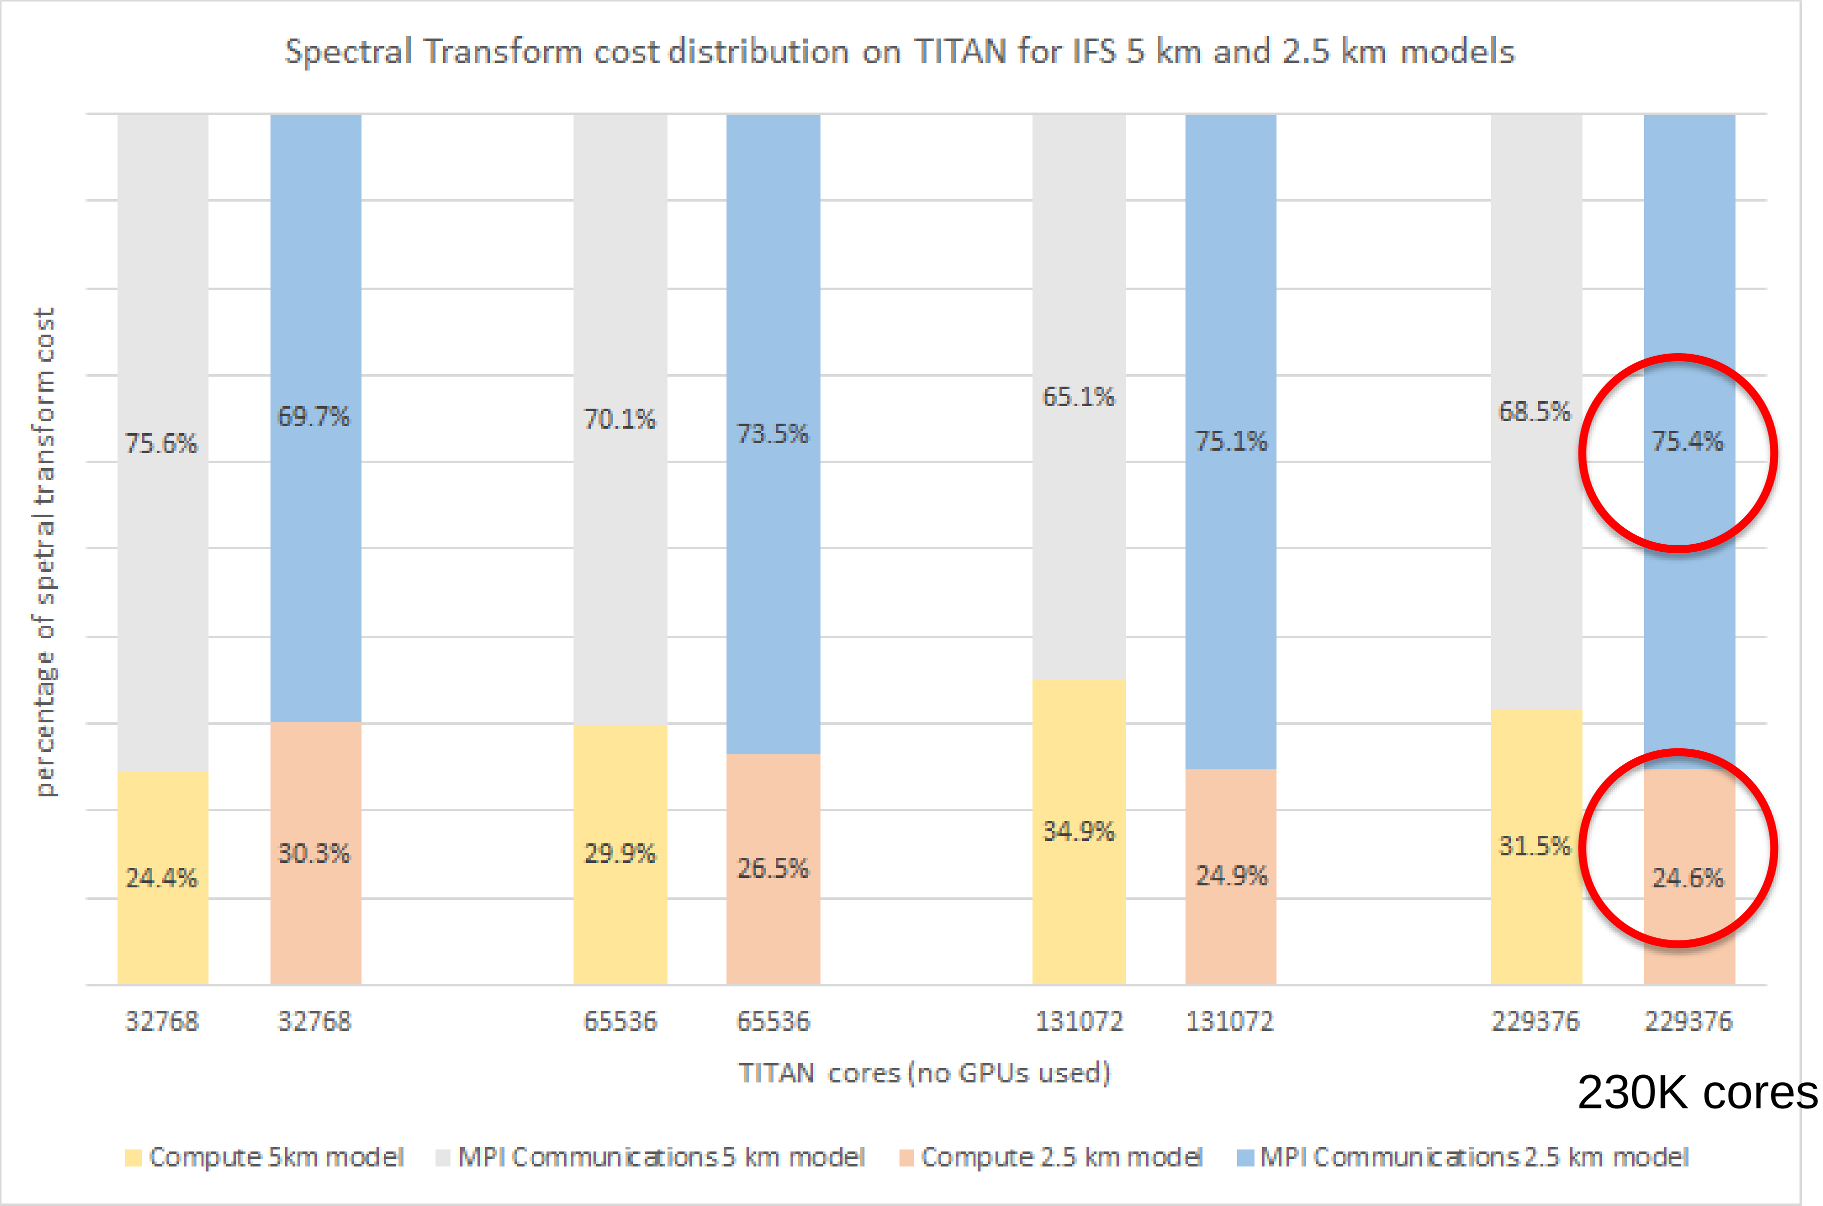

230K cores

Application performance XC30

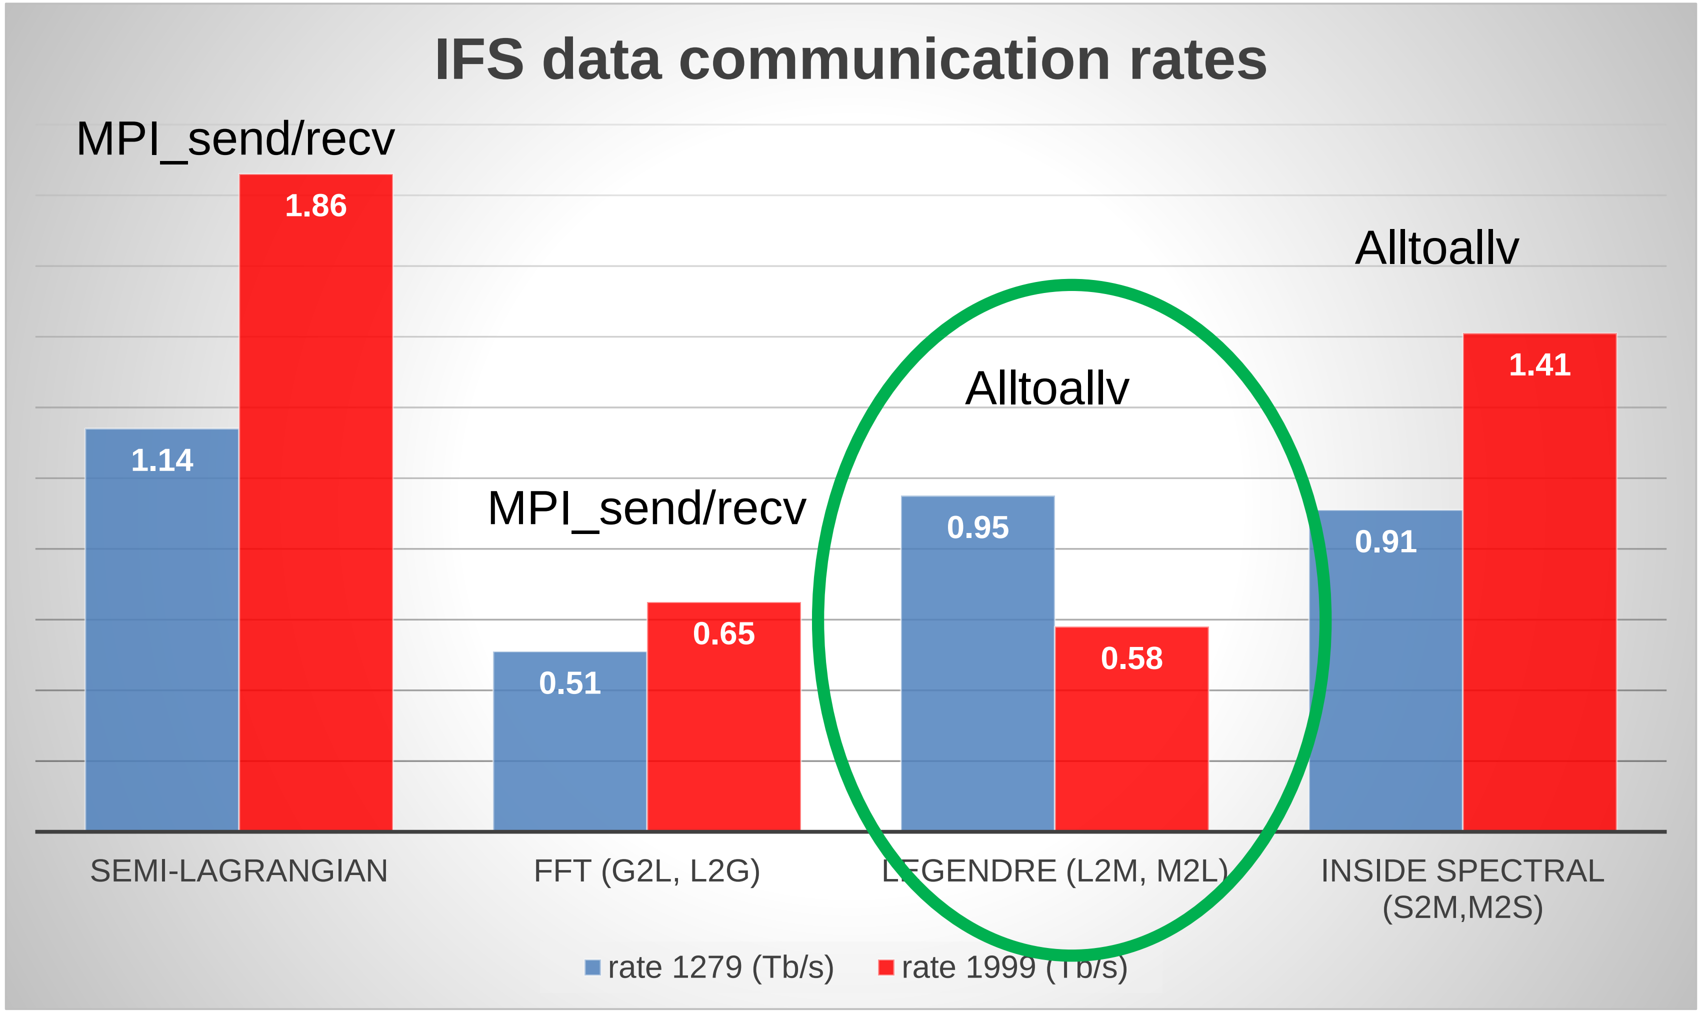

TCo1279 on 360 nodes; TCo1999 on 720 nodes ; dt=450s; 48h forecast

Aries network sketch (at ECMWF)

- On-node == 2 x 12 (now 18) cores
  - "Speed of light"
- On-blade == 4 nodes + switch router with 48 ports using PCIe-3 network cards
  - 15 Gbyte/s
- On-chassis (rank-1) == 16 x blades == using the 48 ports of the router
  - 15 Gbyte/s
- Electrical group (rank-2) == 6 x chassis == 384 nodes
  - 15 Gbyte/s
- Global group (rank-3) == fibre-optical connections ~3500 nodes

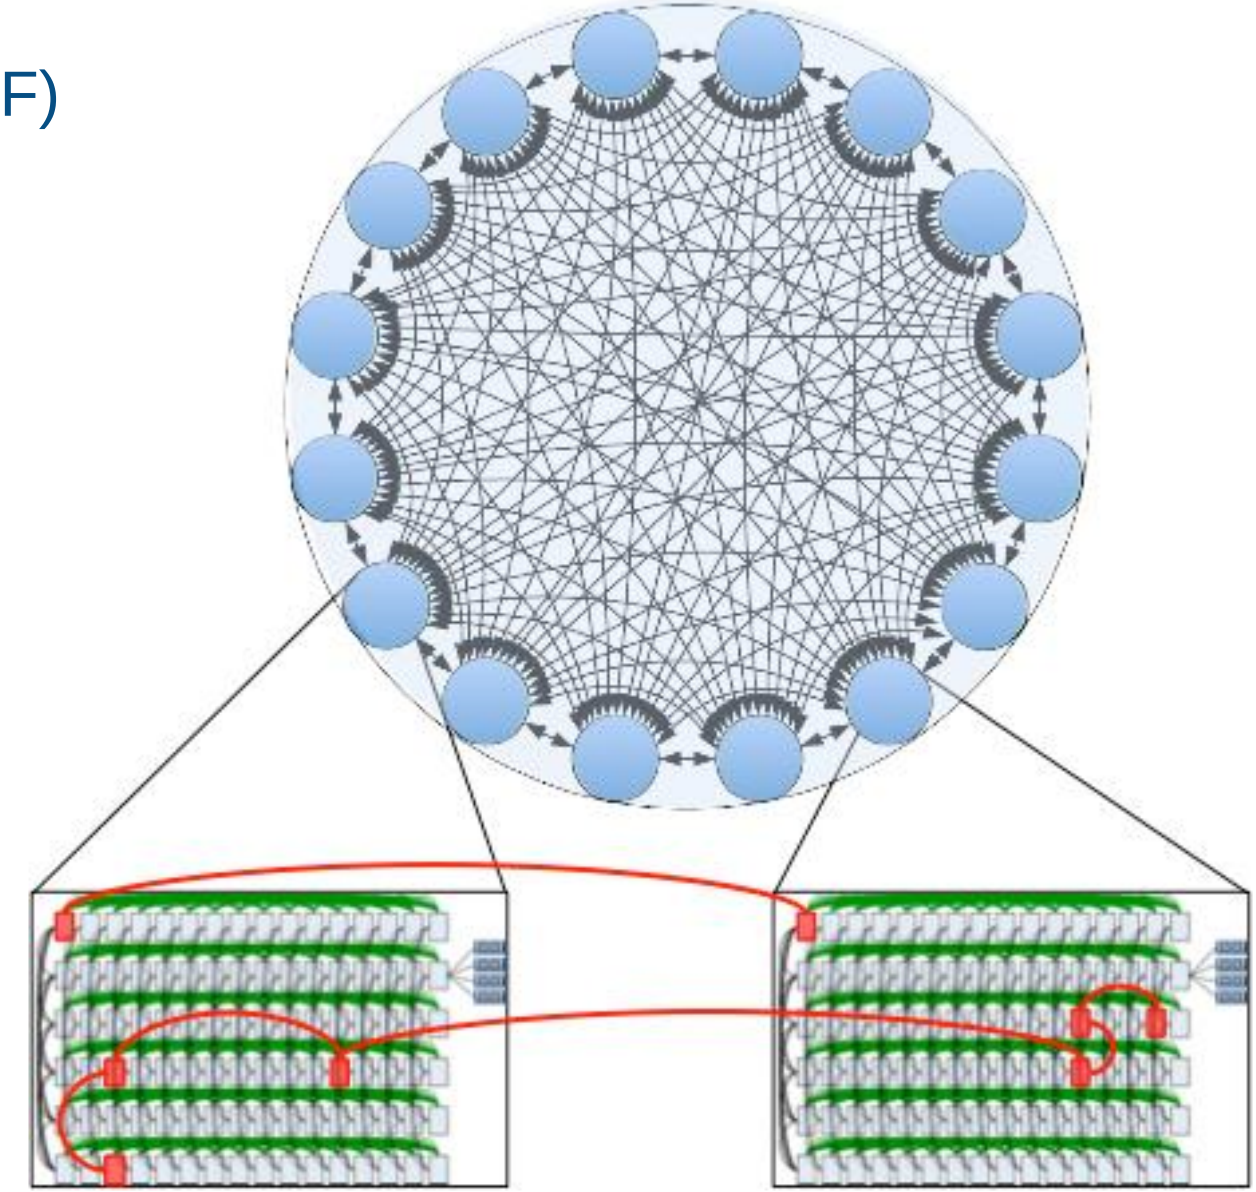

from Brian Austin
NERSC Advanced Technology Group
NUG 2014, February 6, 2014

Comparison to alternative methods

Time-to-solution as a function of CPU cores at ~ 3km

Operational need!

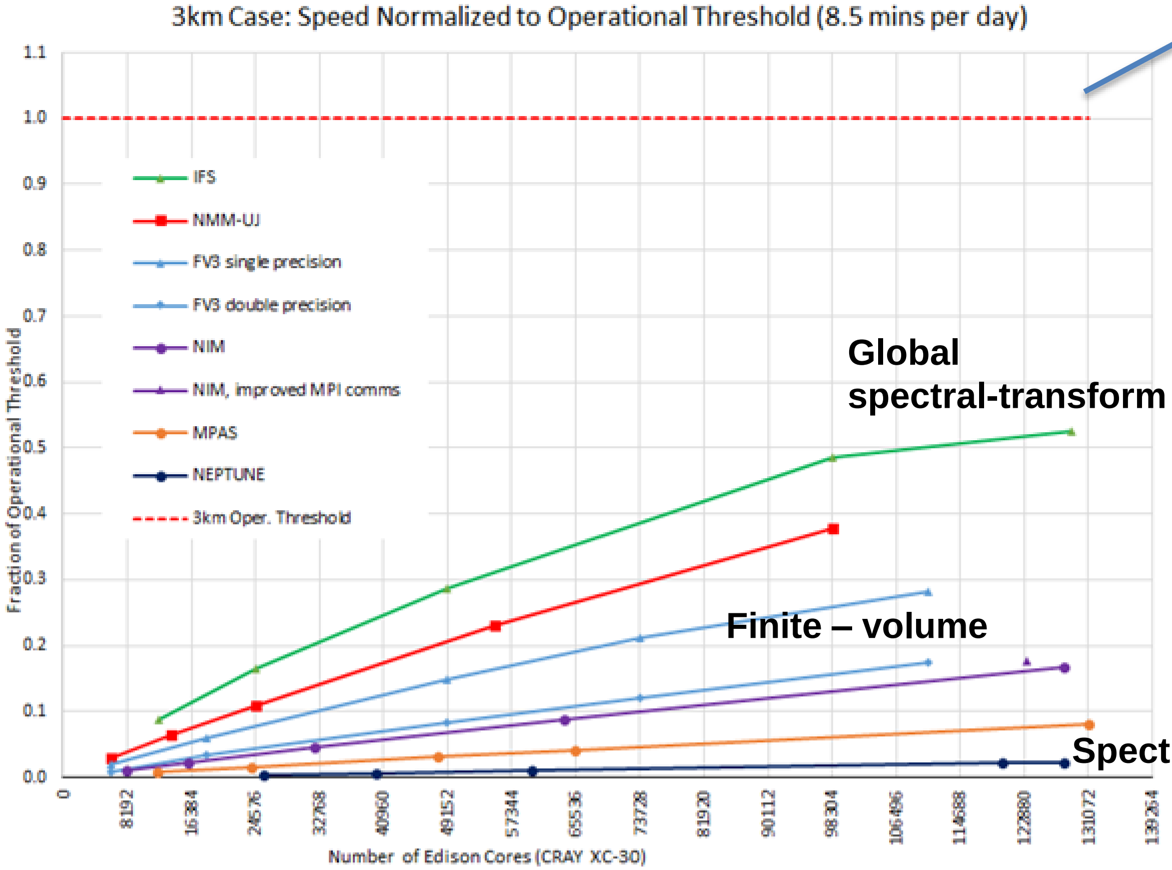

Global spectral-transform - IFS

Finite – volume

Spectral - element

(adapted from Michalakes et al, NGGPS report, 2015)

Projected communication volumes at ~5km resolution at comparable time-to-solution

10 prognostic variables
137 levels

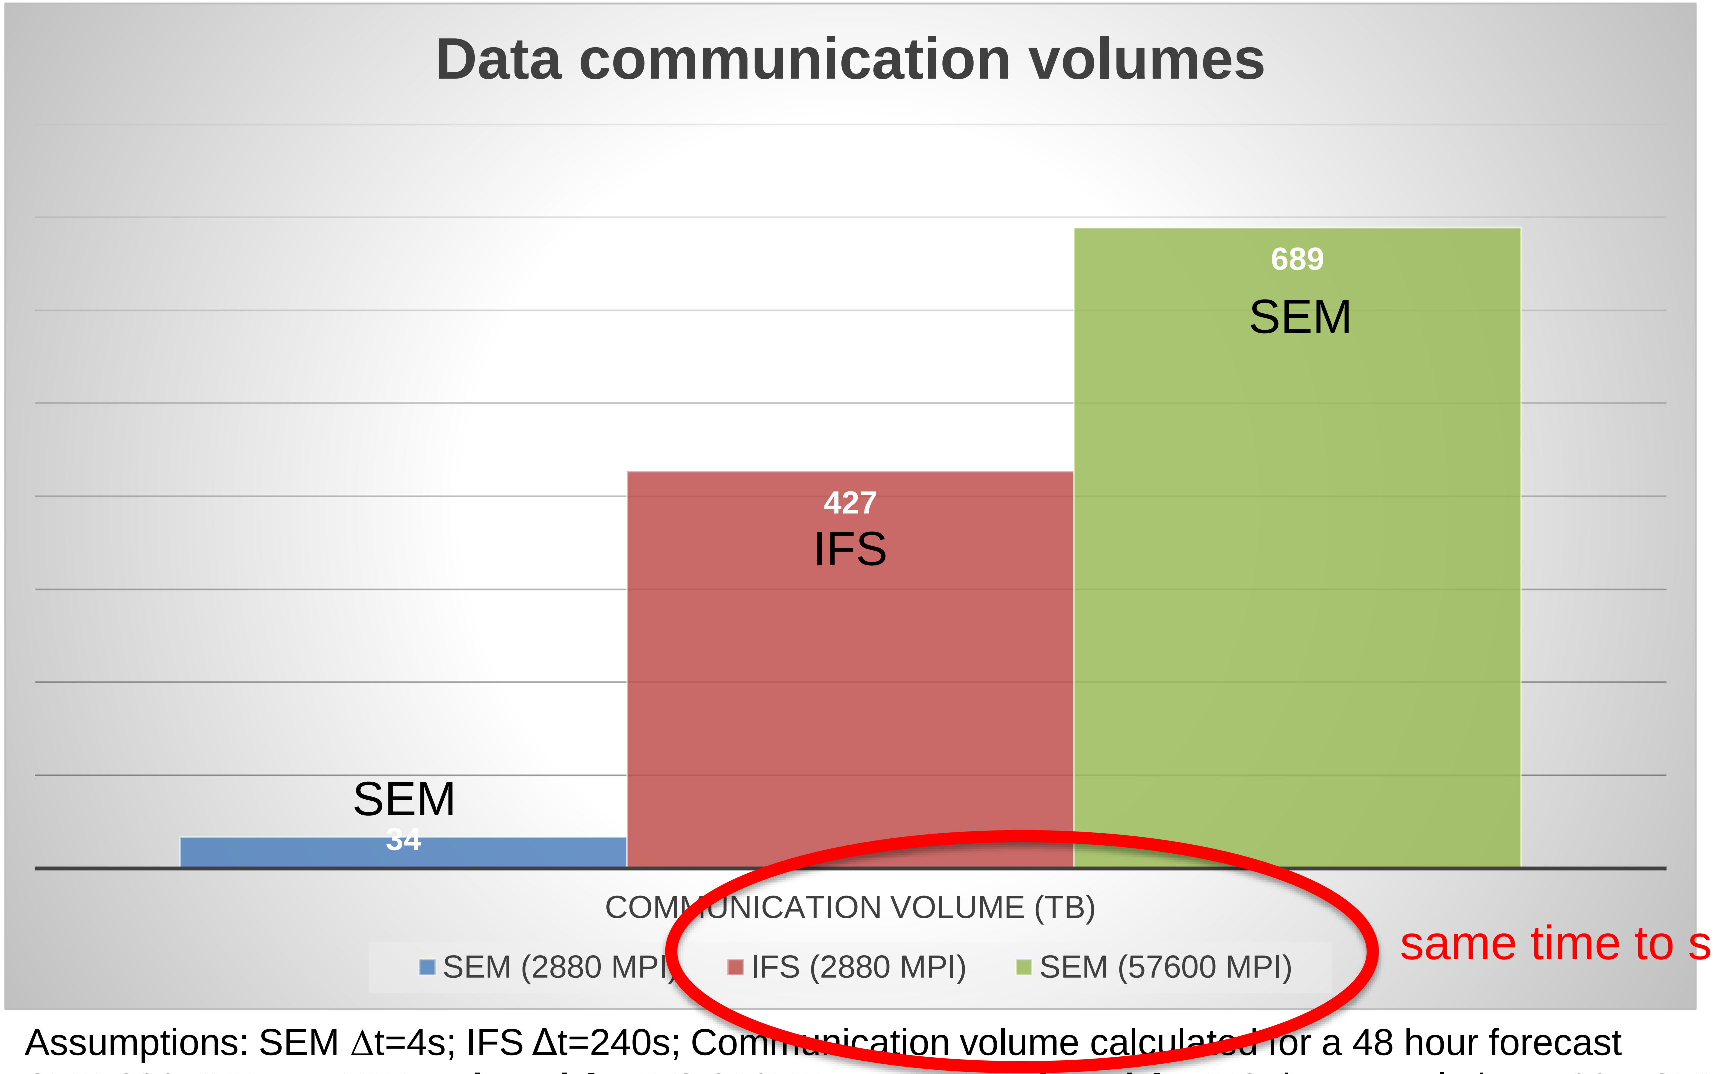

same time to solution

Energy efficiency ?

Emerging HPC
architectures may
change the game ?
>> see A. Mueller talk

Assumptions: SEM Δt=4s; IFS Δt=240s; Communication volume calculated for a 48 hour forecast
SEM 290.4KB per MPI task and Δt; IFS 216MB per MPI task and Δt; IFS time-to-solution = 20 x SEM (Michalakes et al 2015)

Algorithmic flexibility

- Global (pseudo-)spectral transform
- Local low-order (<=2) methods (Finite-volume)
- Local higher-order (>2) methods (spectral element, DG, flux reconstruction)
  - Can these be efficiently maintained in the same numerical modelling framework ?
  - A grid-choice that facilitates hybrid use of several of above simultaneously ?
  - Storage layer abstraction for emerging HPC architectures ?

⇒ *Atlas* framework

(Deconinck et al, 2016)

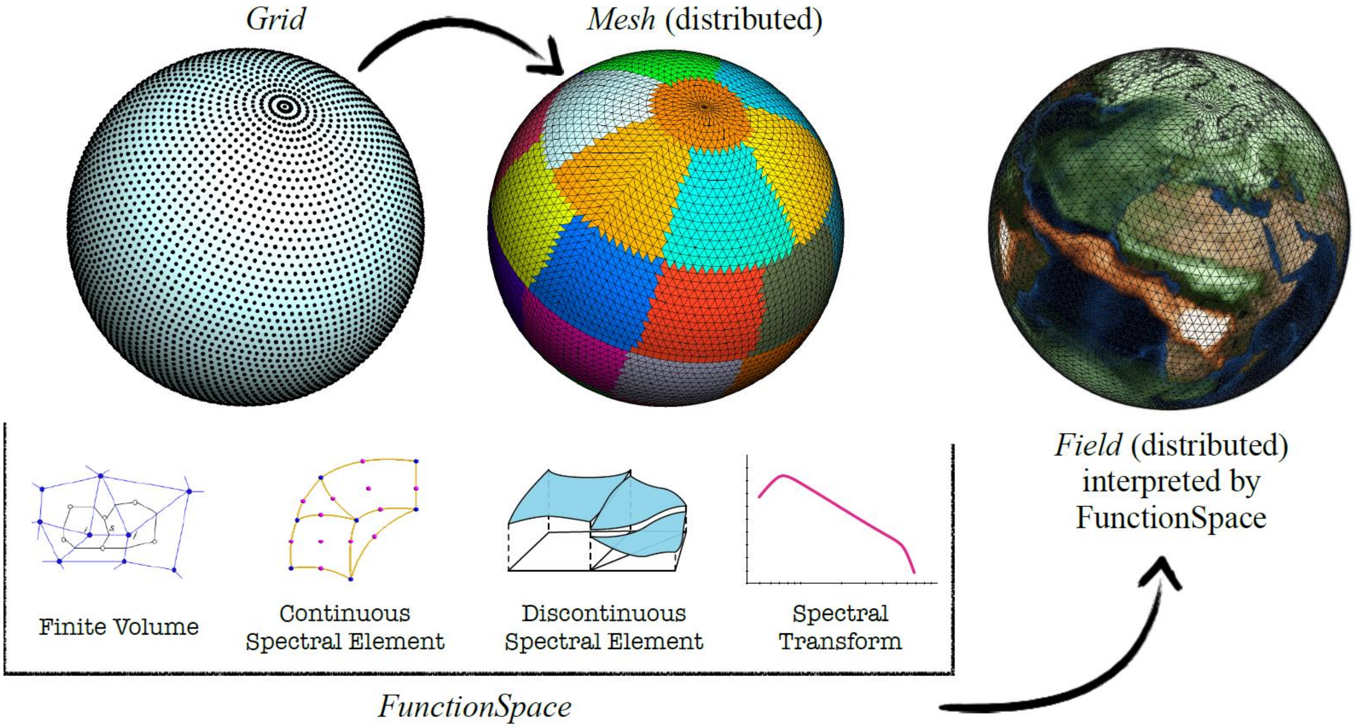

ESCAPE

Funded by the European Union

Energy efficient **SC**alable **A**lgorithms for weather **P**rediction at **E**xascale

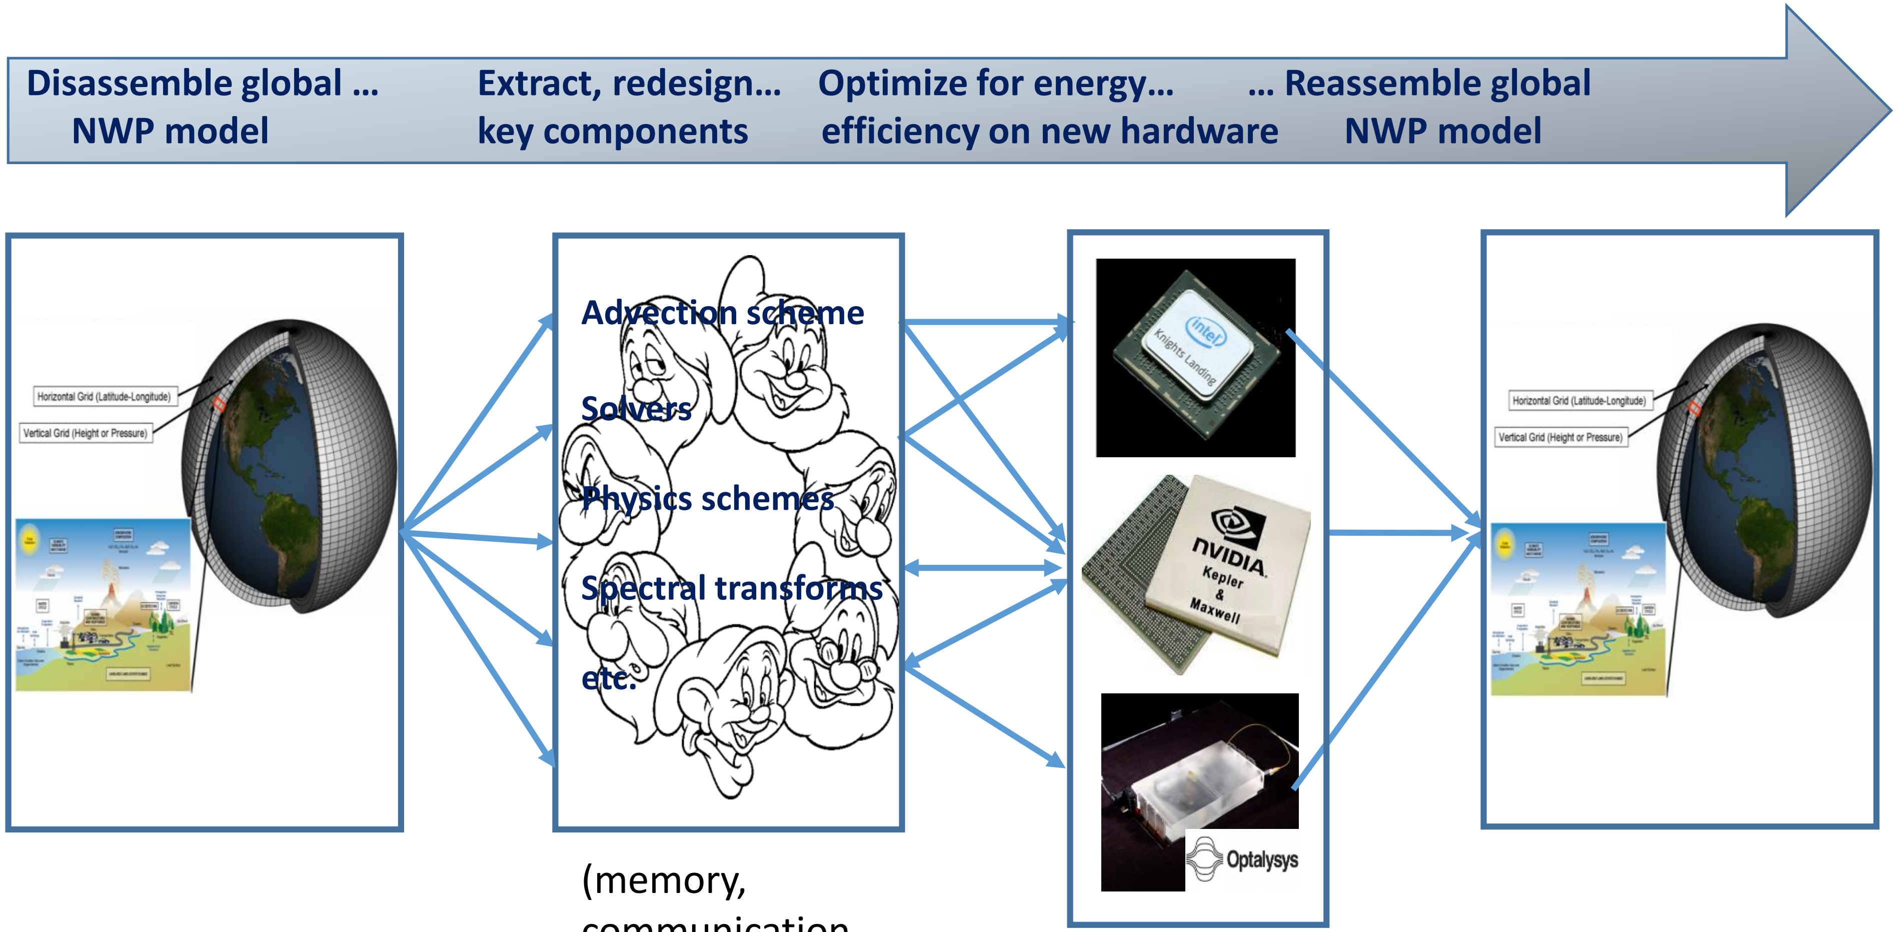

Disassemble global … NWP model

Extract, redesign… key components

Optimize for energy… efficiency on new hardware

… Reassemble global NWP model

- Advection scheme
- Solvers
- Physics schemes
- Spectral transforms
- etc.

(memory, communication, compute intensity)

ECMWF MeteoSwiss RMI NVIDIA METEO FRANCE Toujours un temps d'avance ICHEC Deutscher Wetterdienst Wetter und Klima aus einer Hand DWD Optalysys Bull PSNC Loughborough University

Exploring different schemes

Funded by the European Union

Rossby-Haurwitz test case after 7 days

Path-based

dwarf-D-advection-SemiLagrangian

Control-volume-based

dwarf-D-advection-MPDATA

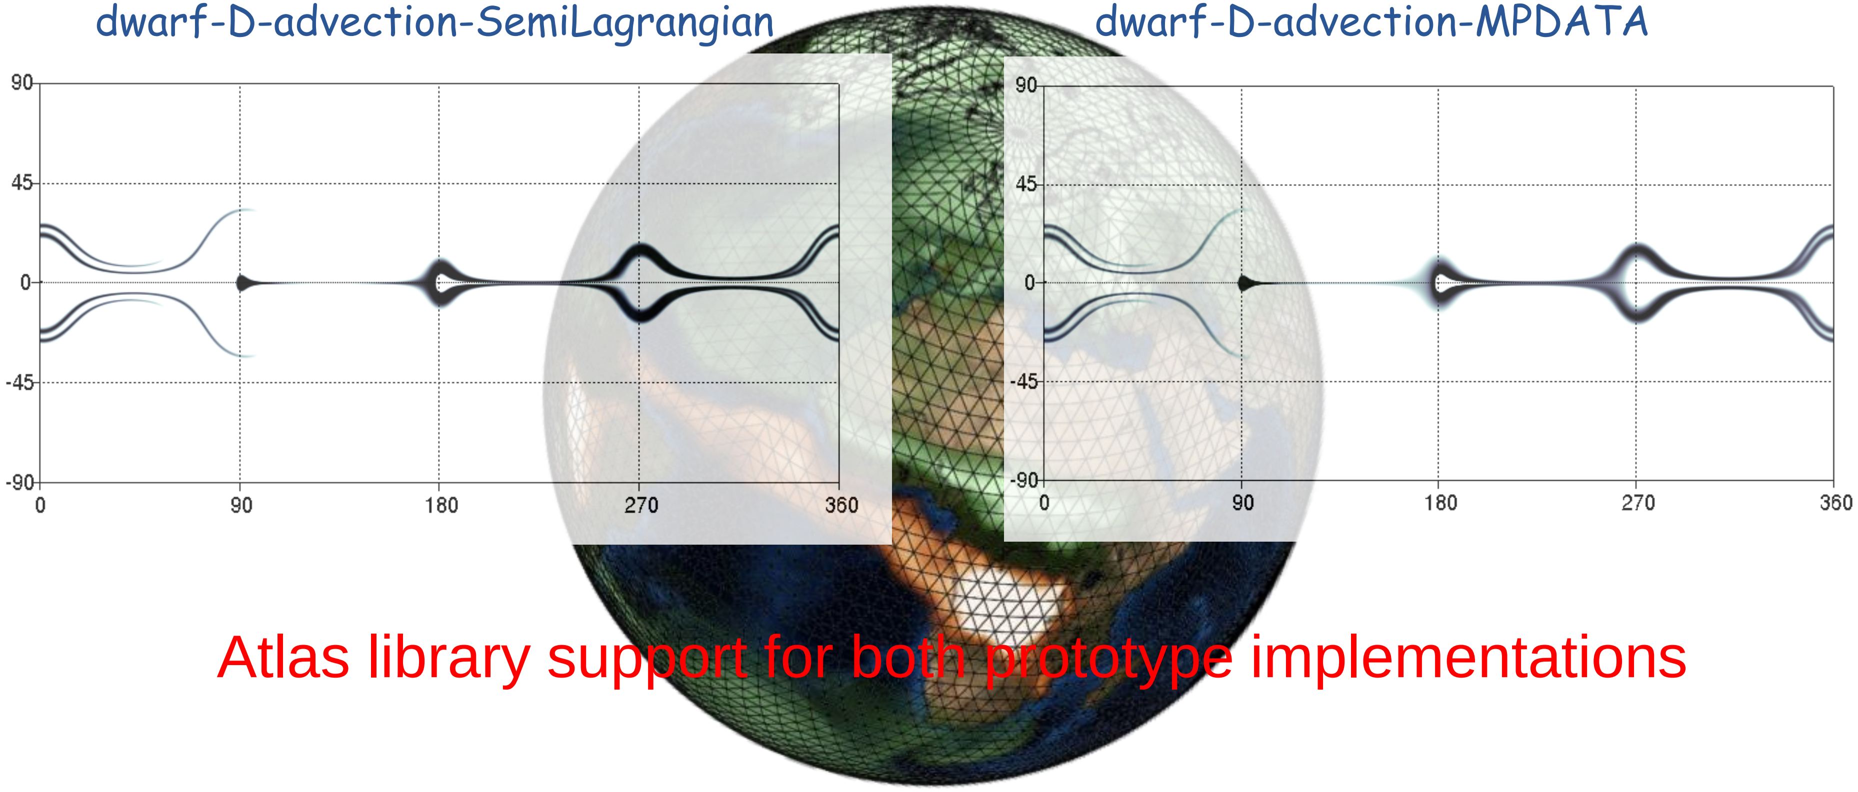

Atlas library support for both prototype implementations

Dry baroclinic instability, FVM (O640) versus the spectral IFS ($T_{co}639$):

Discretization and solvers:

Spectral transform IFS compared to finite-volume FVM

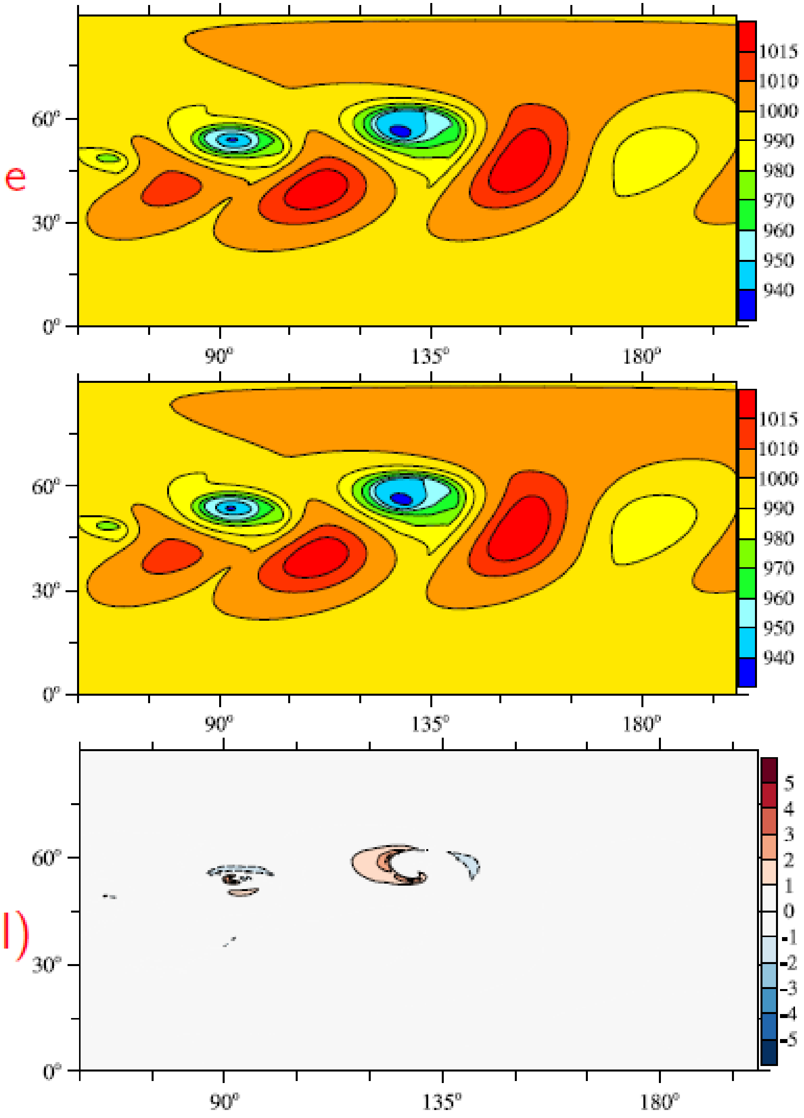

Finite-volume

Spectral

day 10

Difference (FV-Spectral)

Surface pressure (hPa)

European Research Council
Established by the European Commission
Supporting top researchers from anywhere in the world

Funded by the
European Union

NWP & Climate Dwarfs in ESCAPE

Dwarf		Target Accelerators			Programming Models			
	Priority	GPU	Xeon Phi	Optalysis	MPI	OpenMP	OpenACC	DSL
D - spectral transform - SH	I	✔	✔	✘	✔	✔	✏	✘
D - spectral transform - biFFT	I	✔	✔	✔	✔	✔	✏	✘
D - advection - MPDATA	I	✔	✔	✘	✔	✔	✏	✏
I - LAITRI (3d interpol. algorithm)	I	✔	✔	✘	✘	✔	✏	✏
D - elliptic solver - GCR	II	✔	✔	✘	✔	✔	✏	⟳
D - advection - semi-Lagrangian	II	✔	✔	✘	✔	✔	✏	✏
P - cloud microphysics - CloudSC	II	✔	✔	✘	✘	✔	⟳	✘
P - radiation scheme - ACRANEB2	II	✔	✔	✘	✘	✔	✏	✘

HPC accelerators challenge the respective roles of targeted “low-level” optimization versus “high-level” algorithmic and code design changes

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 671627

Atmospheric composition: high resolution CO_2 forecast

Interacting global atmosphere, land-surface, radiation, ocean, wave, and chemical processes

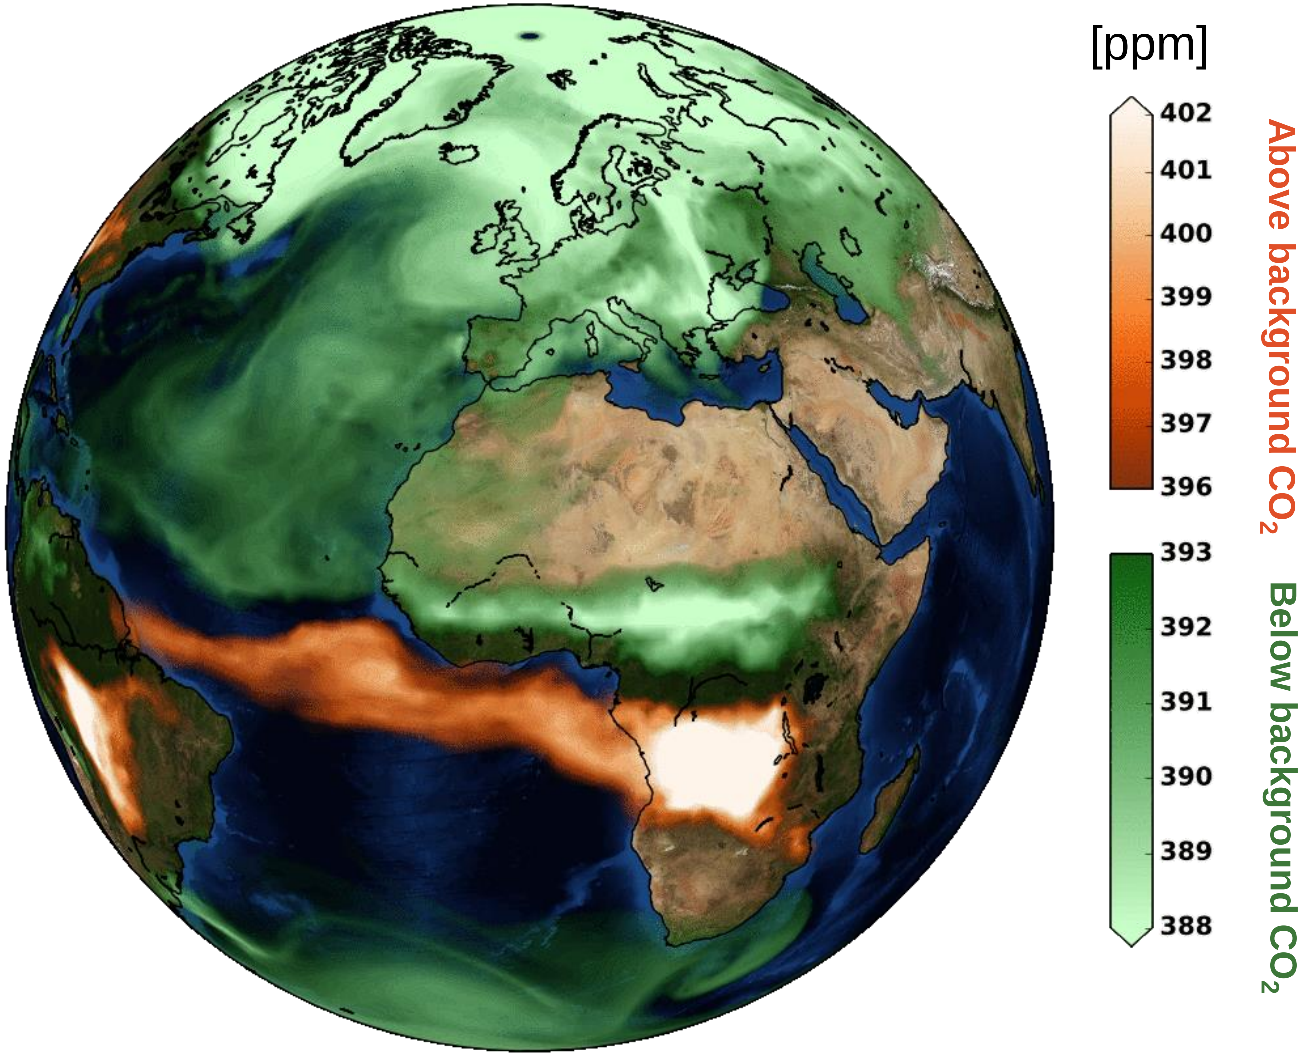

Total column average atmospheric CO_2: September 2013

A. Agusti-Panareda and S. Massart

Cost distribution with increased complexity

TCo1279 L137 (~9km)
coupled to ORCA025_Z75, WAM + LIM2 sea-ice

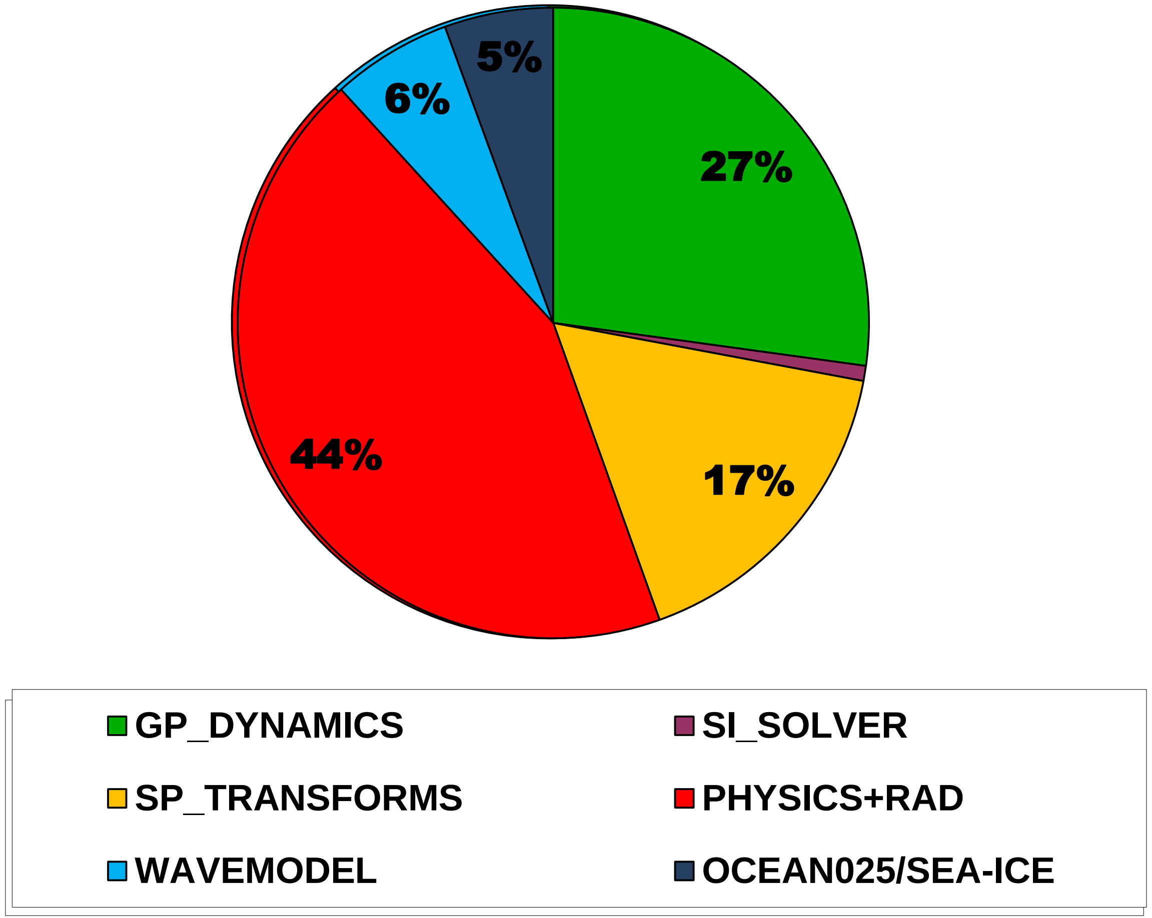

TCo639 L91 (~18km)
coupled to WAM

Thanks to J. Flemming

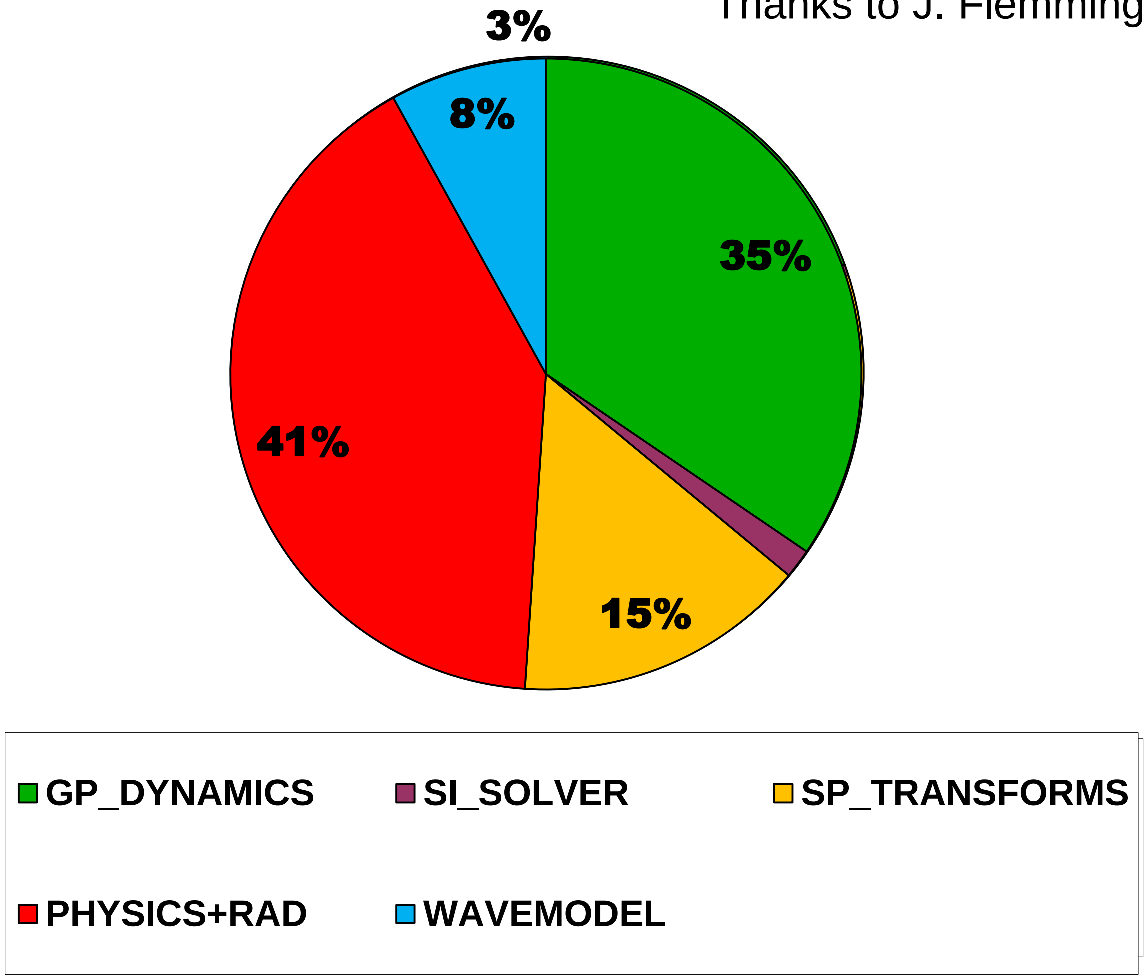

Cost distribution future atmosphere only

Both runs on 960 nodes XC40 with
1920 MPI tasks x 18 threads

TCo1999 L62 (~5km)
10 Billion DOF
~ Memory 3.5Gb / node

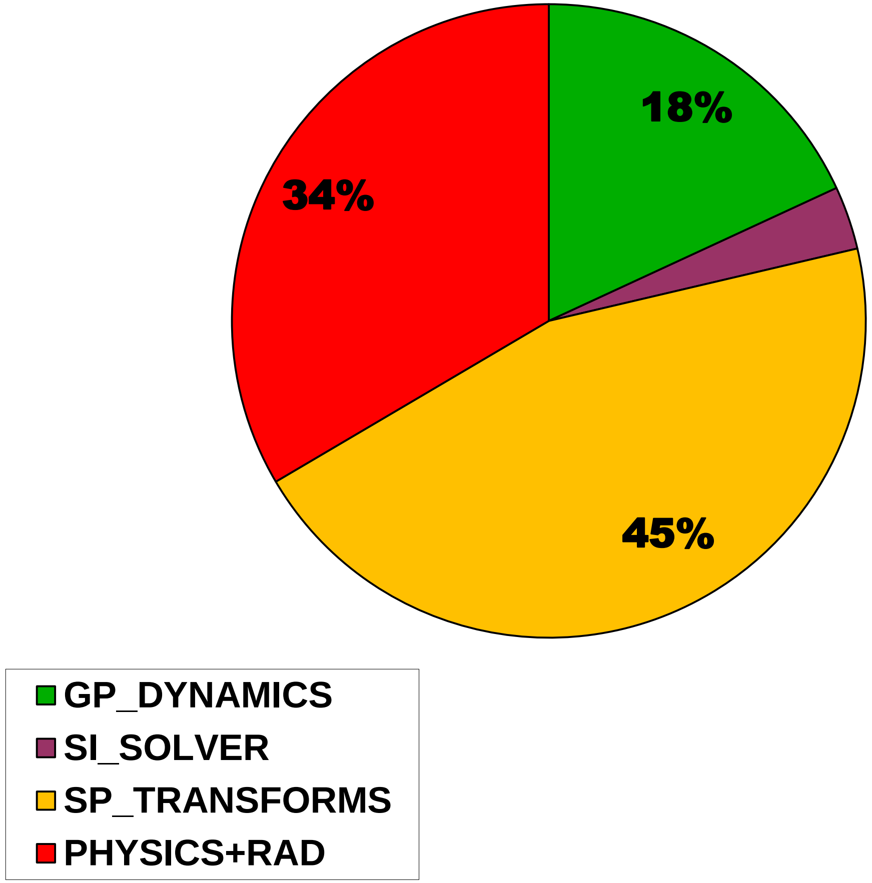

TCo7999 L62 (~1.3km)
160 Billion DOF
~ Memory 35Gb / node

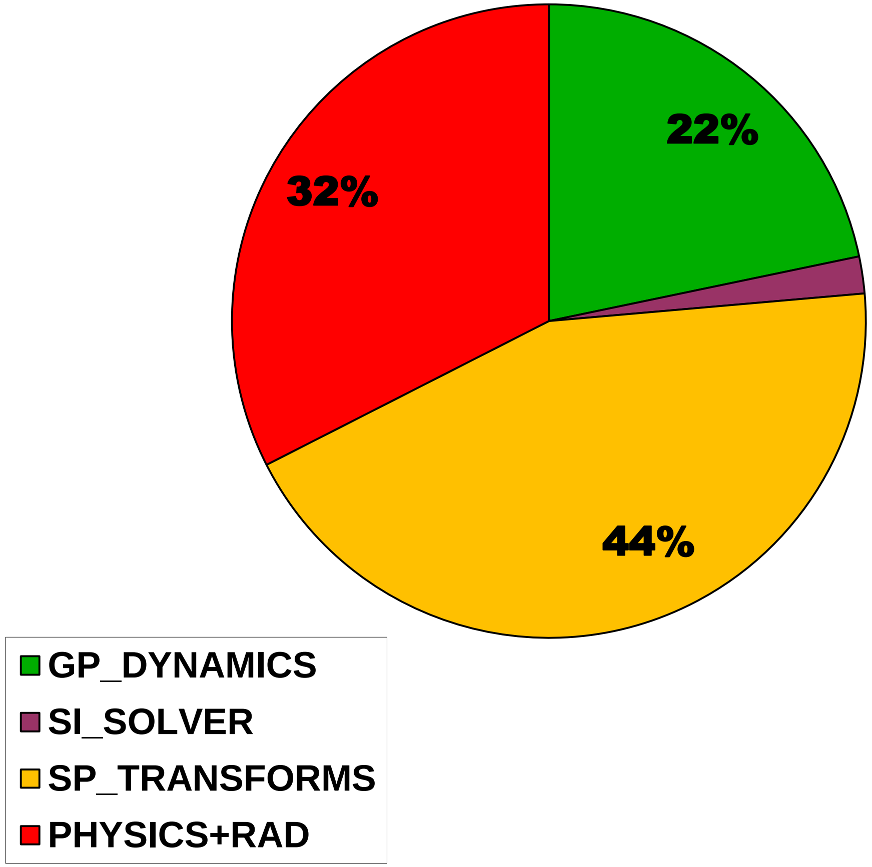

TCo1279 (~9km) World's highest resolution global NWP today

TCo1279 (~9km) a 6.6 Megapixel camera with predictive skill!

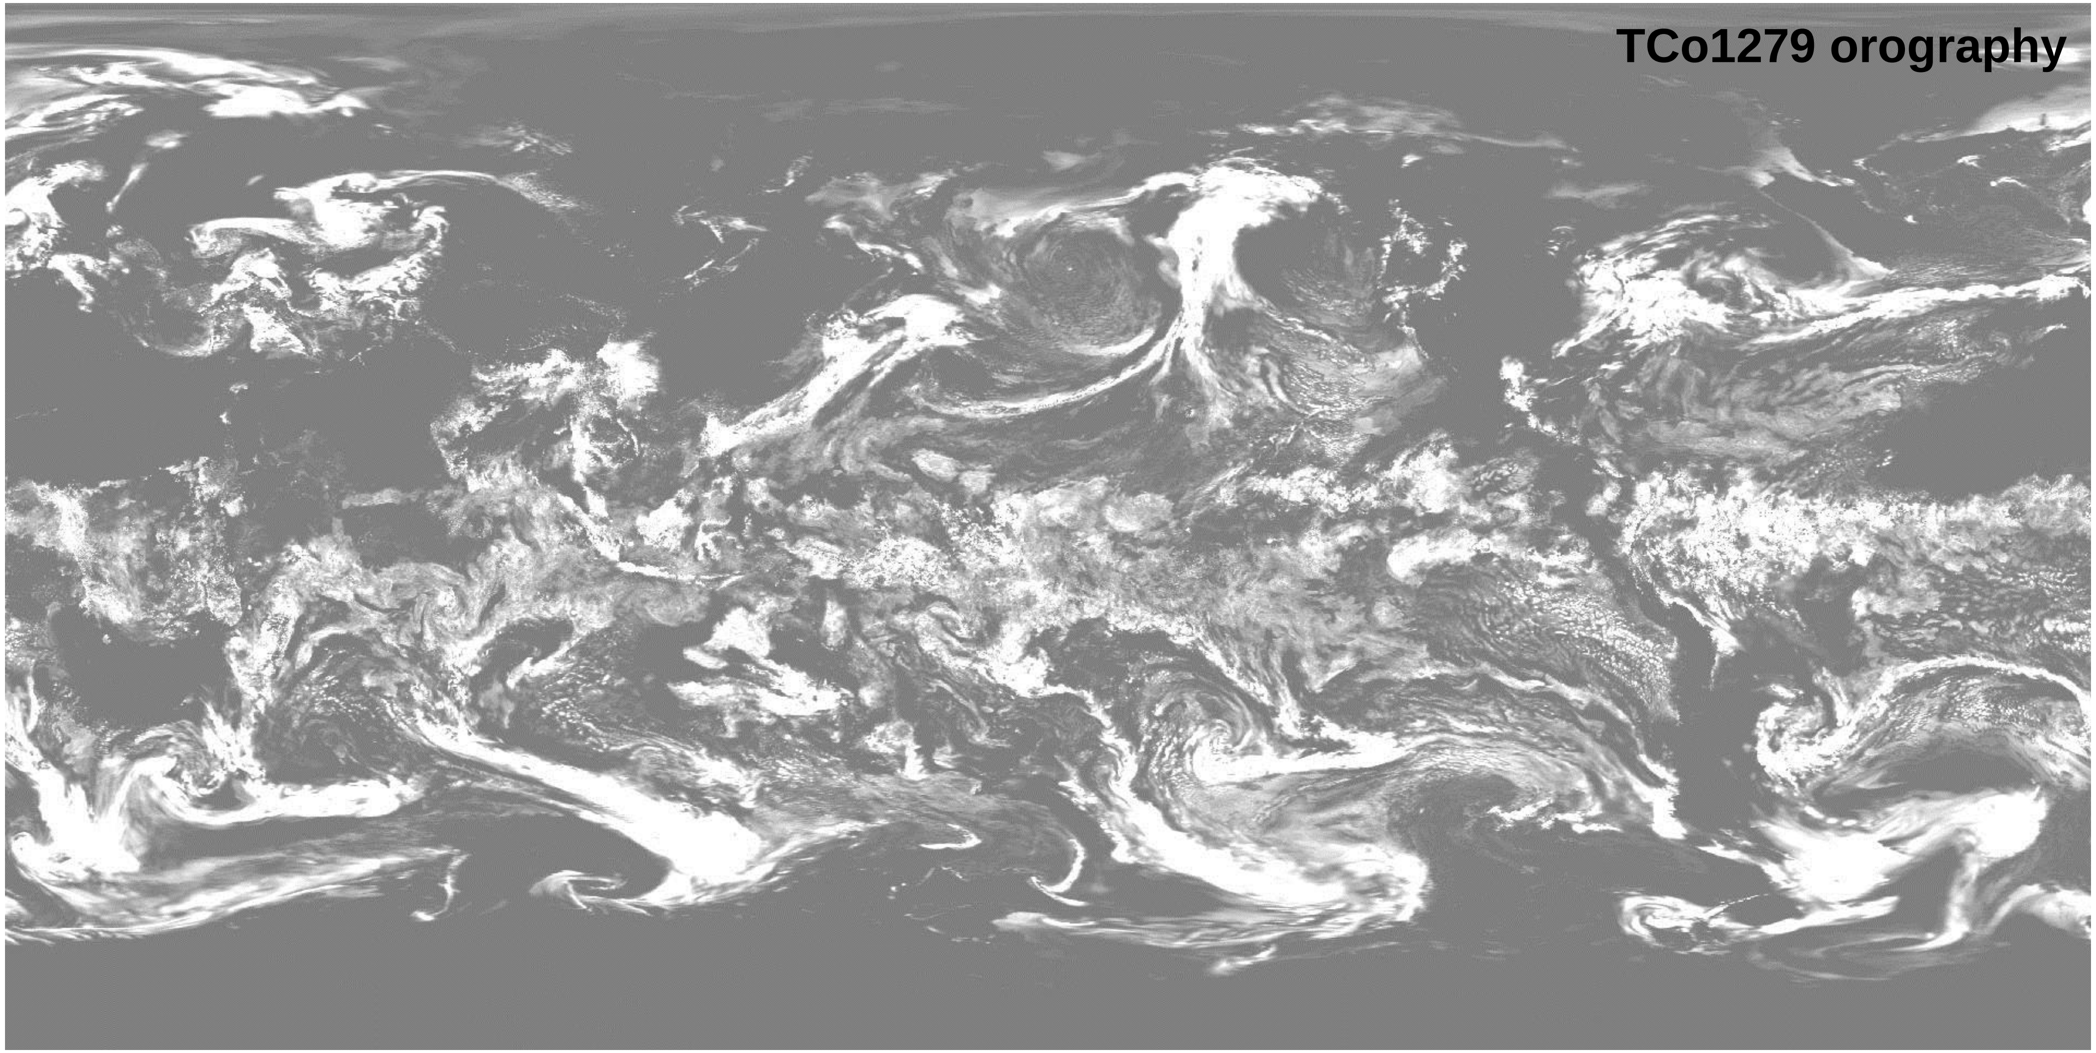

(12h forecast, *hydrostatic*, with deep convection parametrization, 450s time-step, 240 Broadwell nodes, ~0.75s per timestep)

TCo7999 (~1.3km) 256 Megapixel camera with predictive skill!

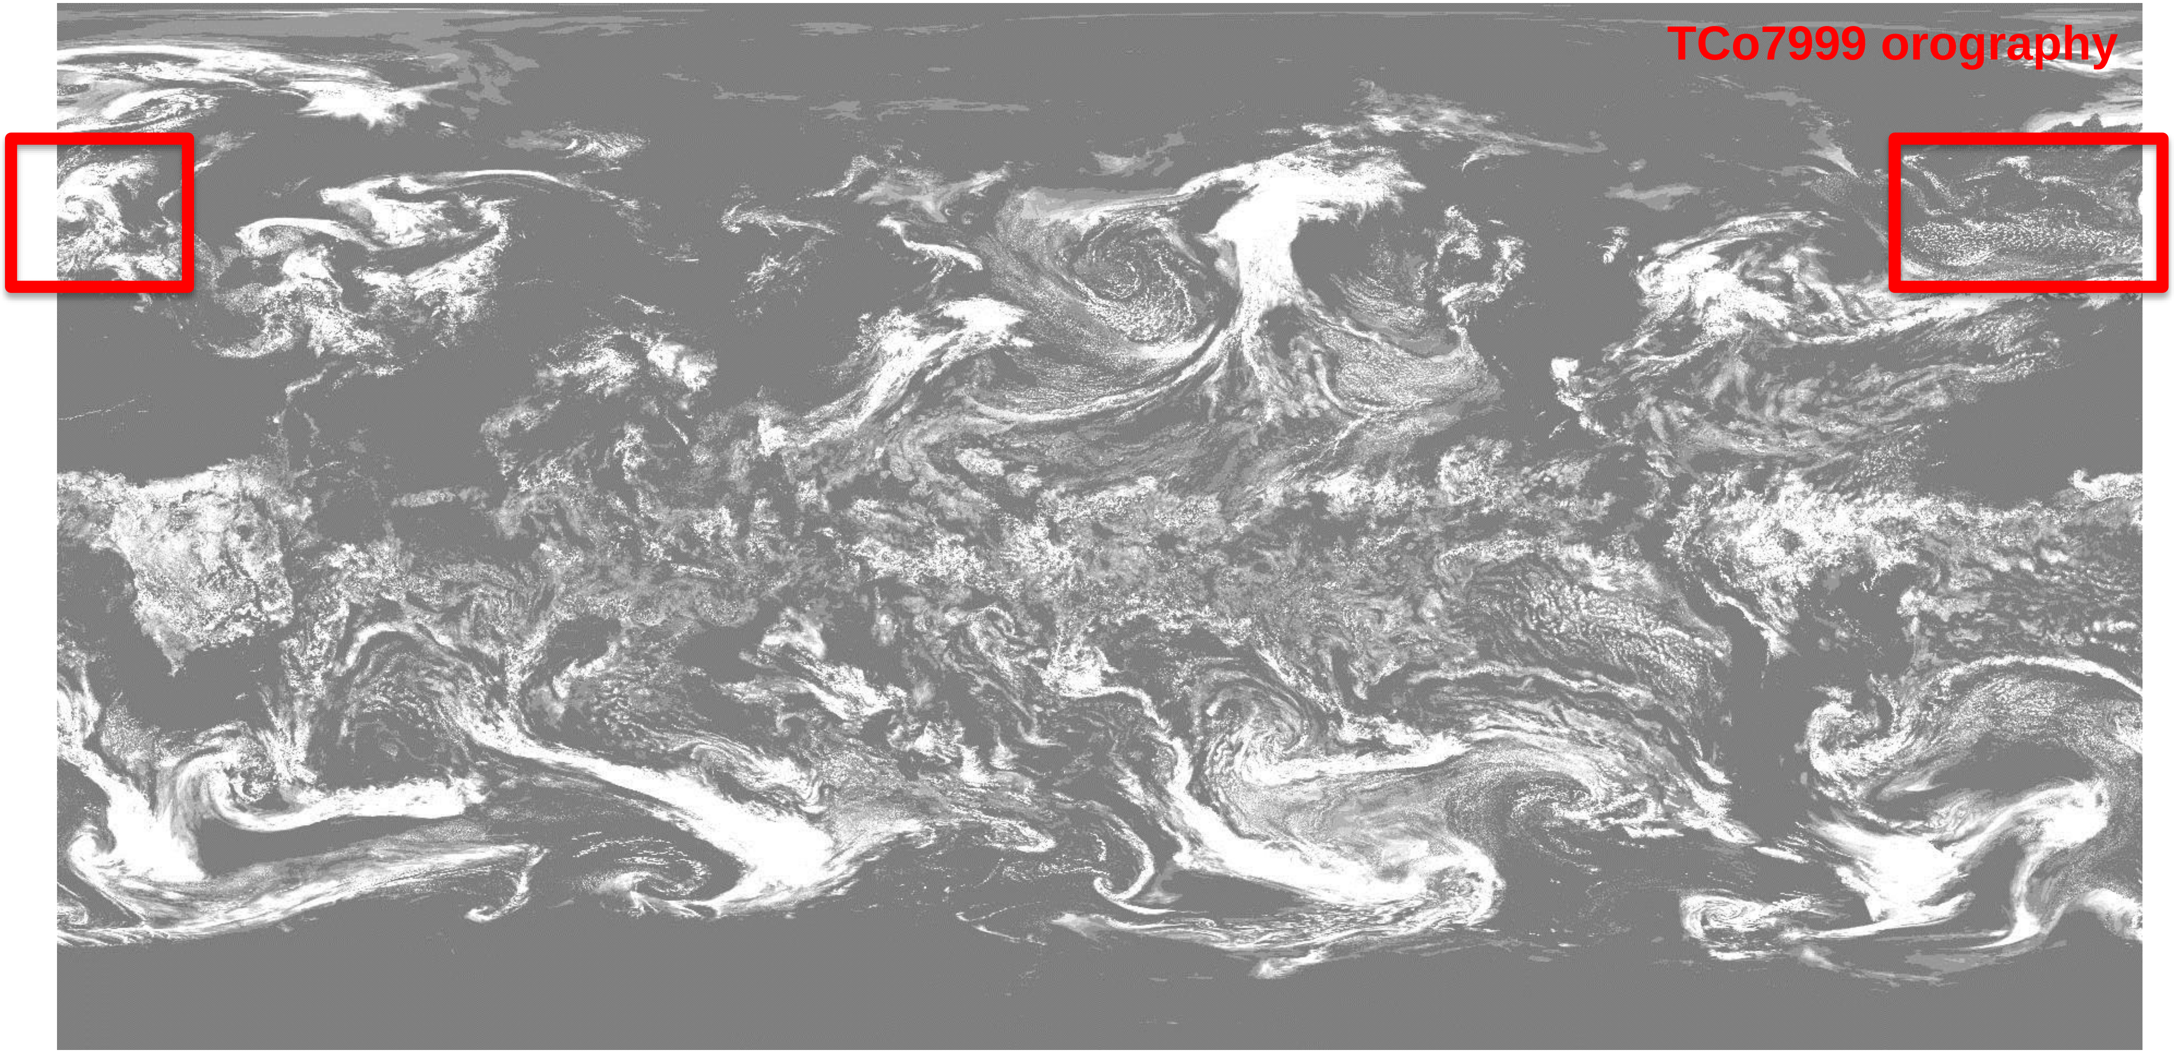

(12 h forecast, *hydrostatic*, no deep convection parametrization, 120s time-step, 960 Broadwell nodes, ~10s per timestep)

Needs zooming …. Windstorm Katie in March 2016 TCo7999 (~1.3km) 03UTC

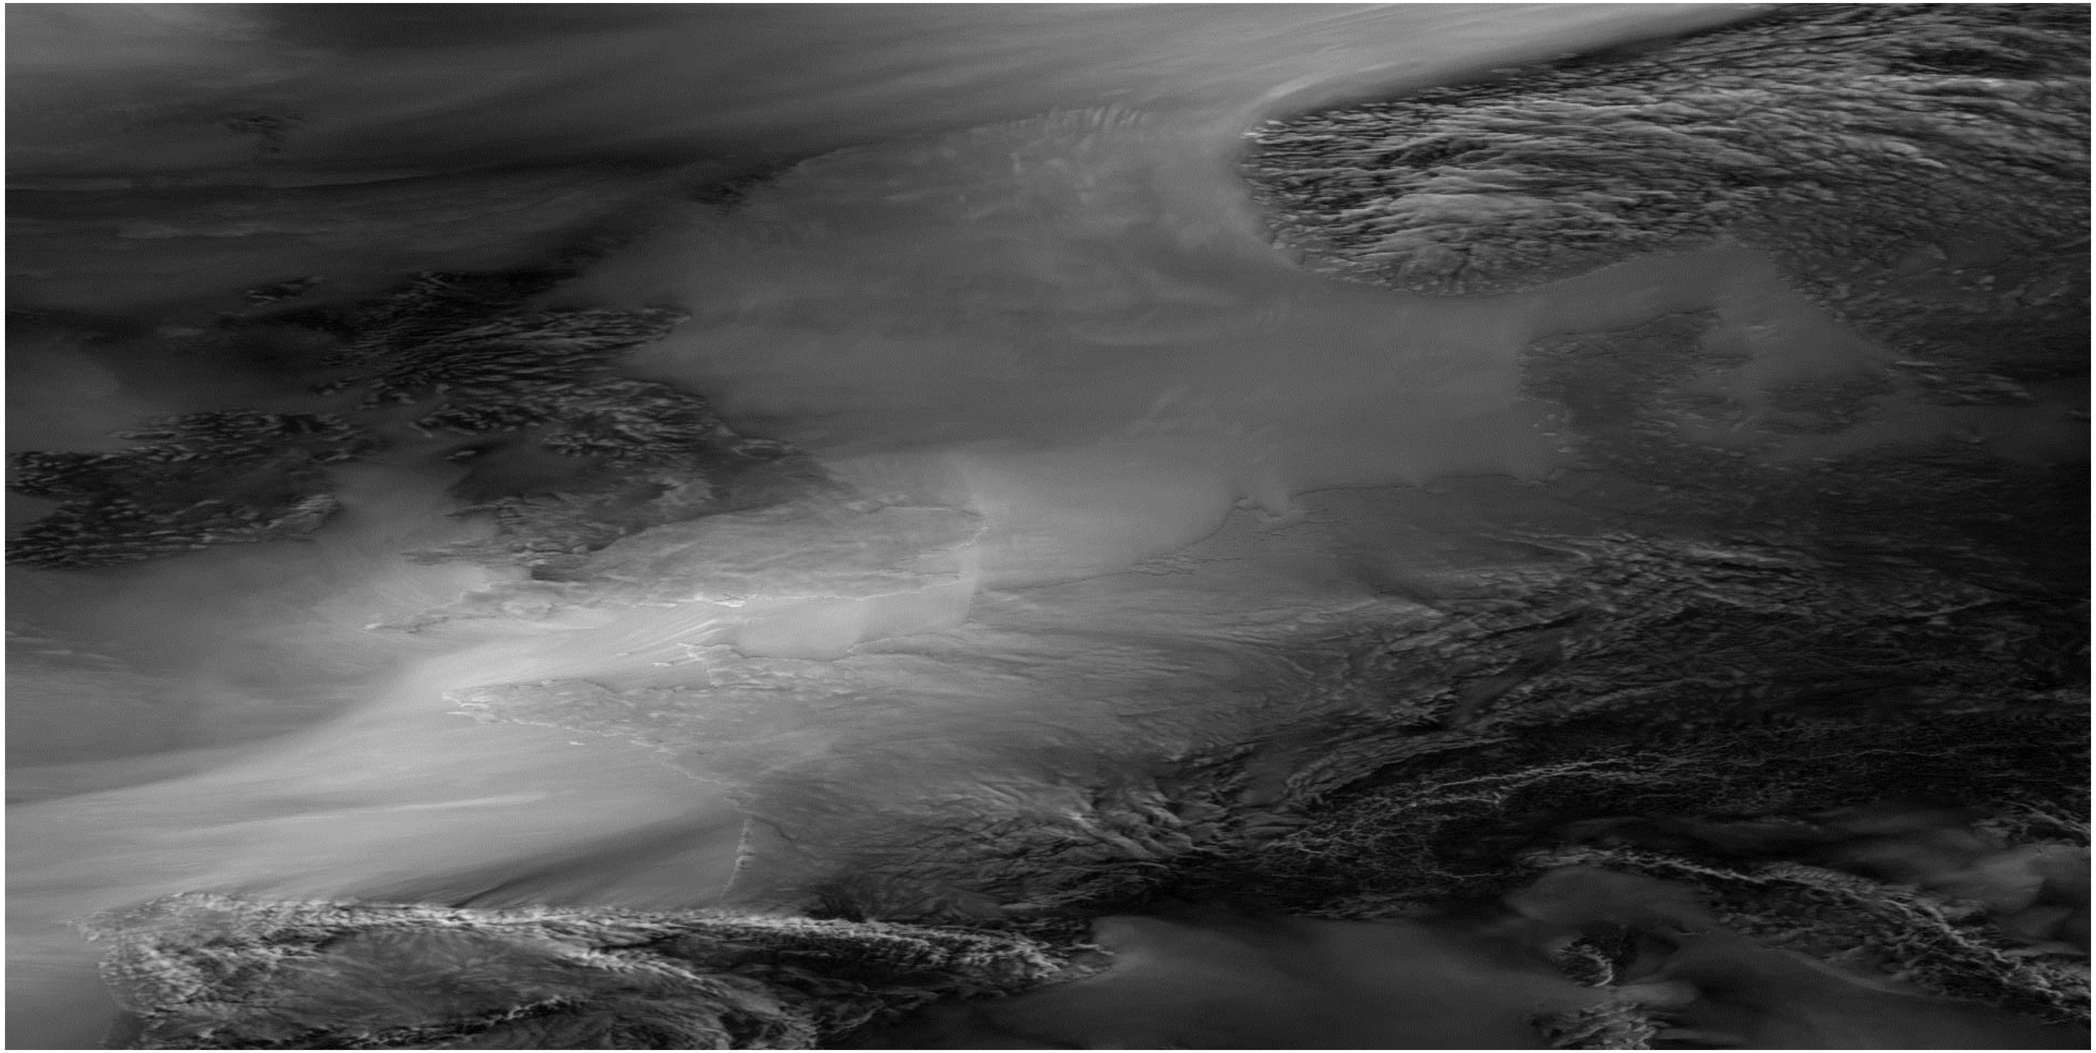

Needs zooming …. Windstorm Katie in March 2016 TCo7999 (~1.3km) 06UTC

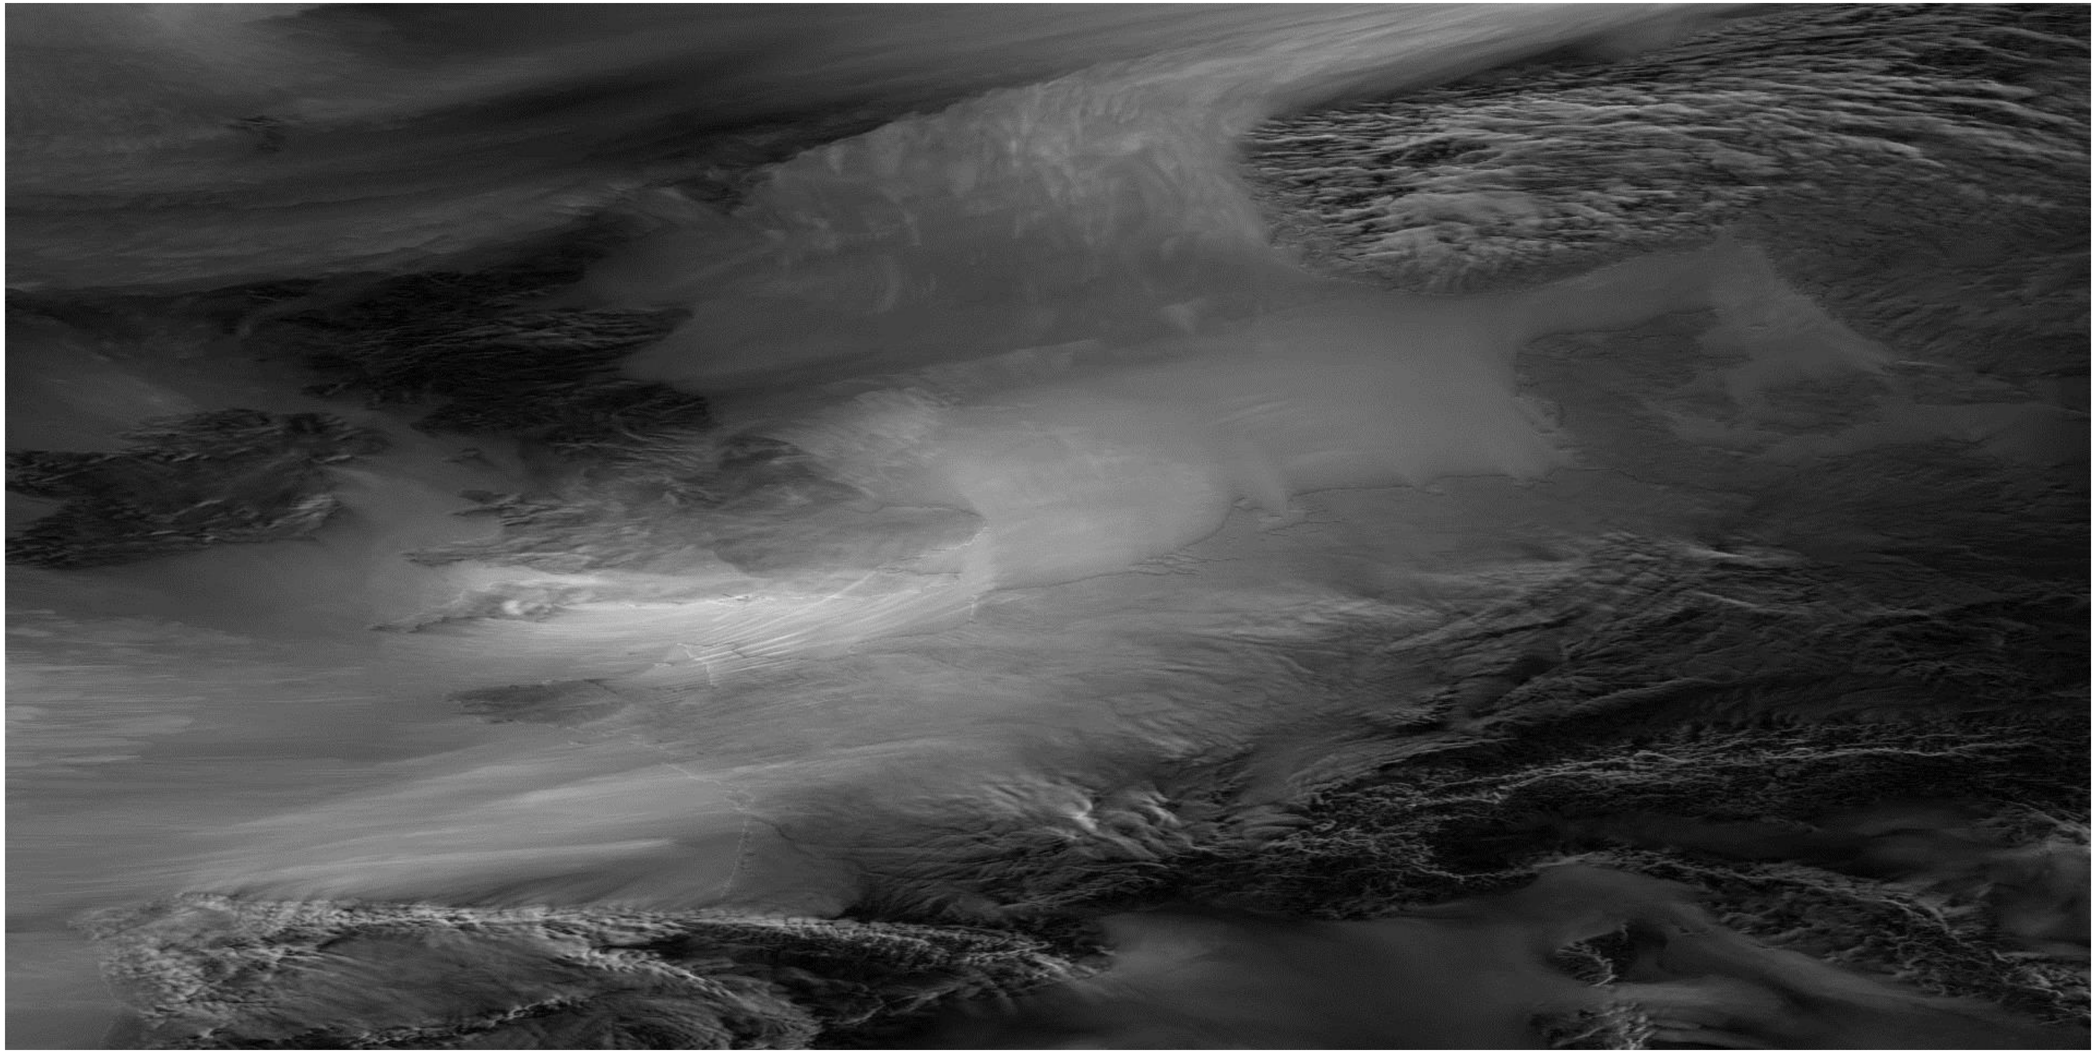

Needs zooming …. Windstorm Katie in March 2016 TCo7999 (~1.3km) 09UTC

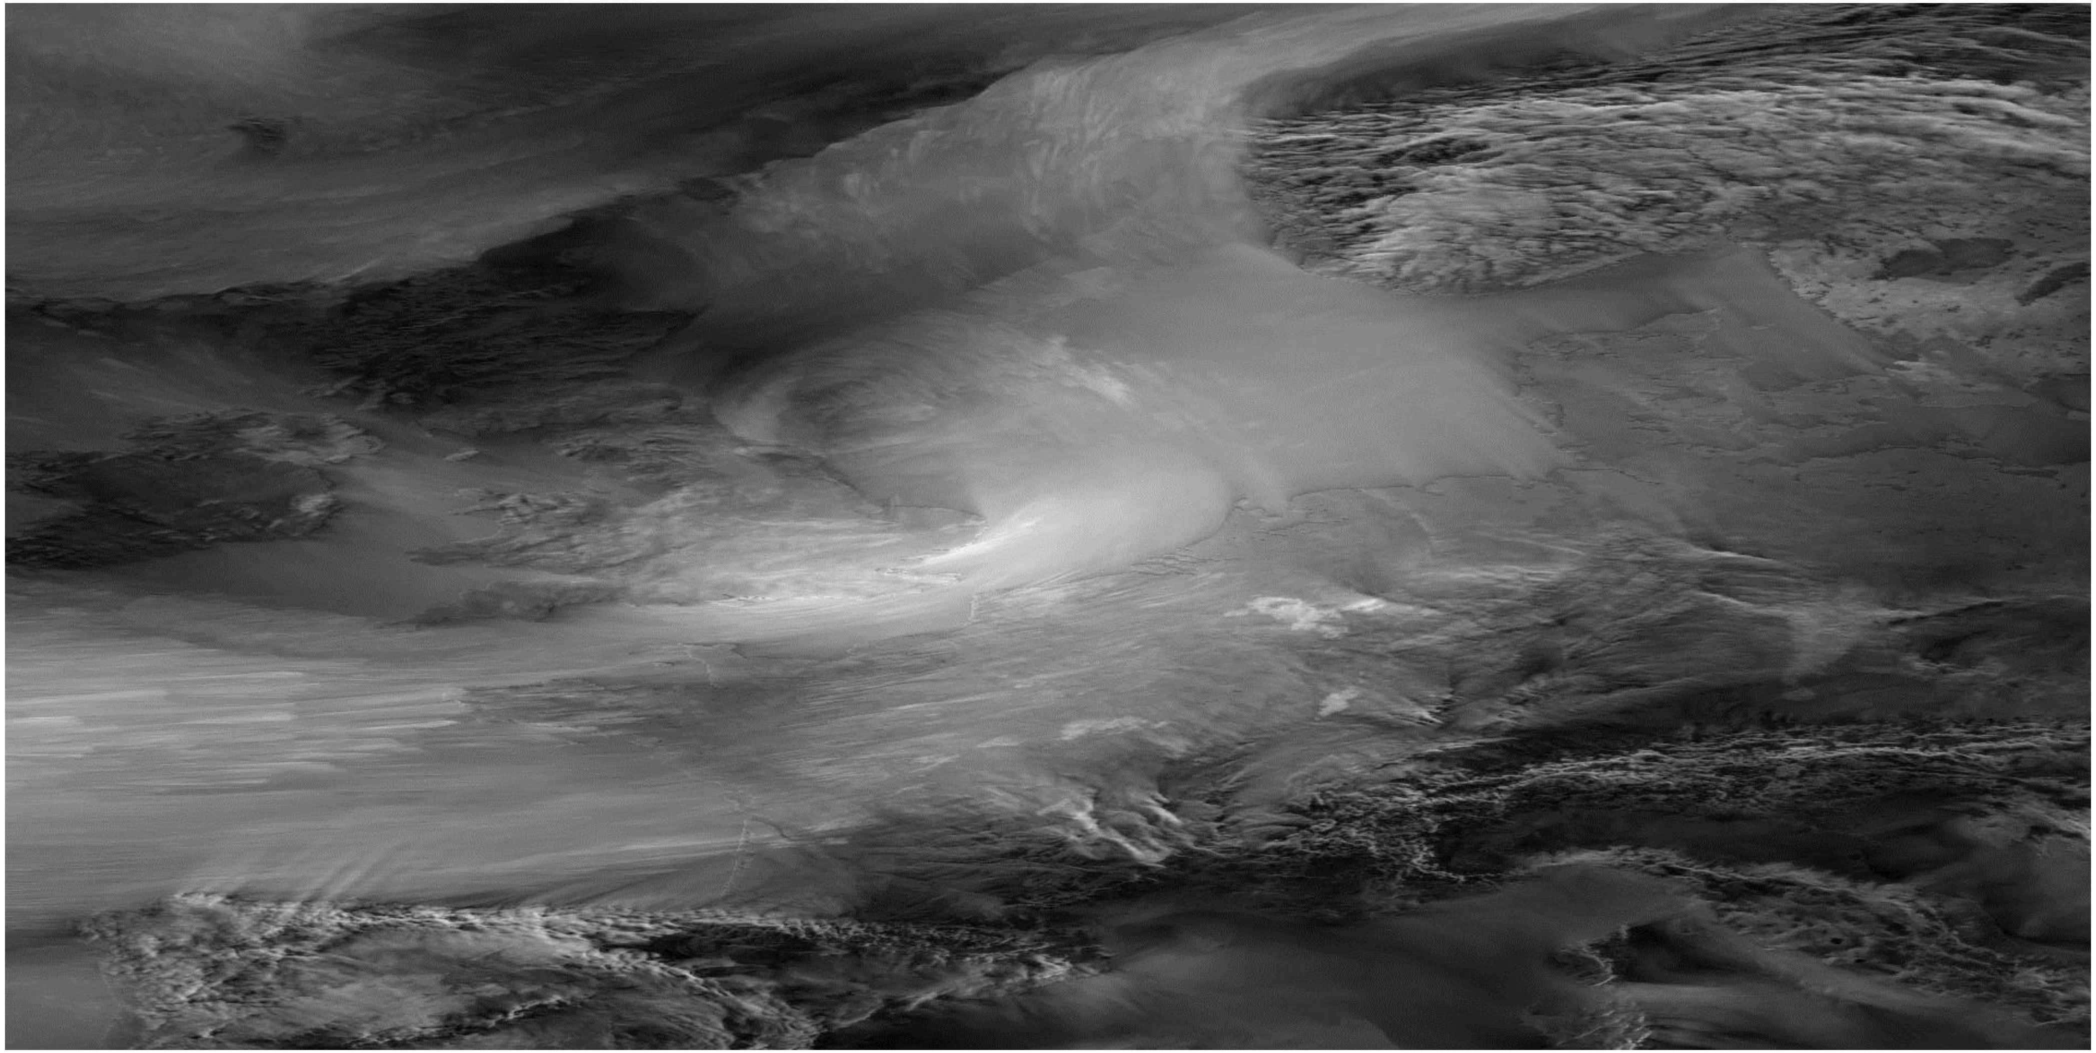

Needs zooming …. Windstorm Katie in March 2016 TCo7999 (~1.3km) 12UTC

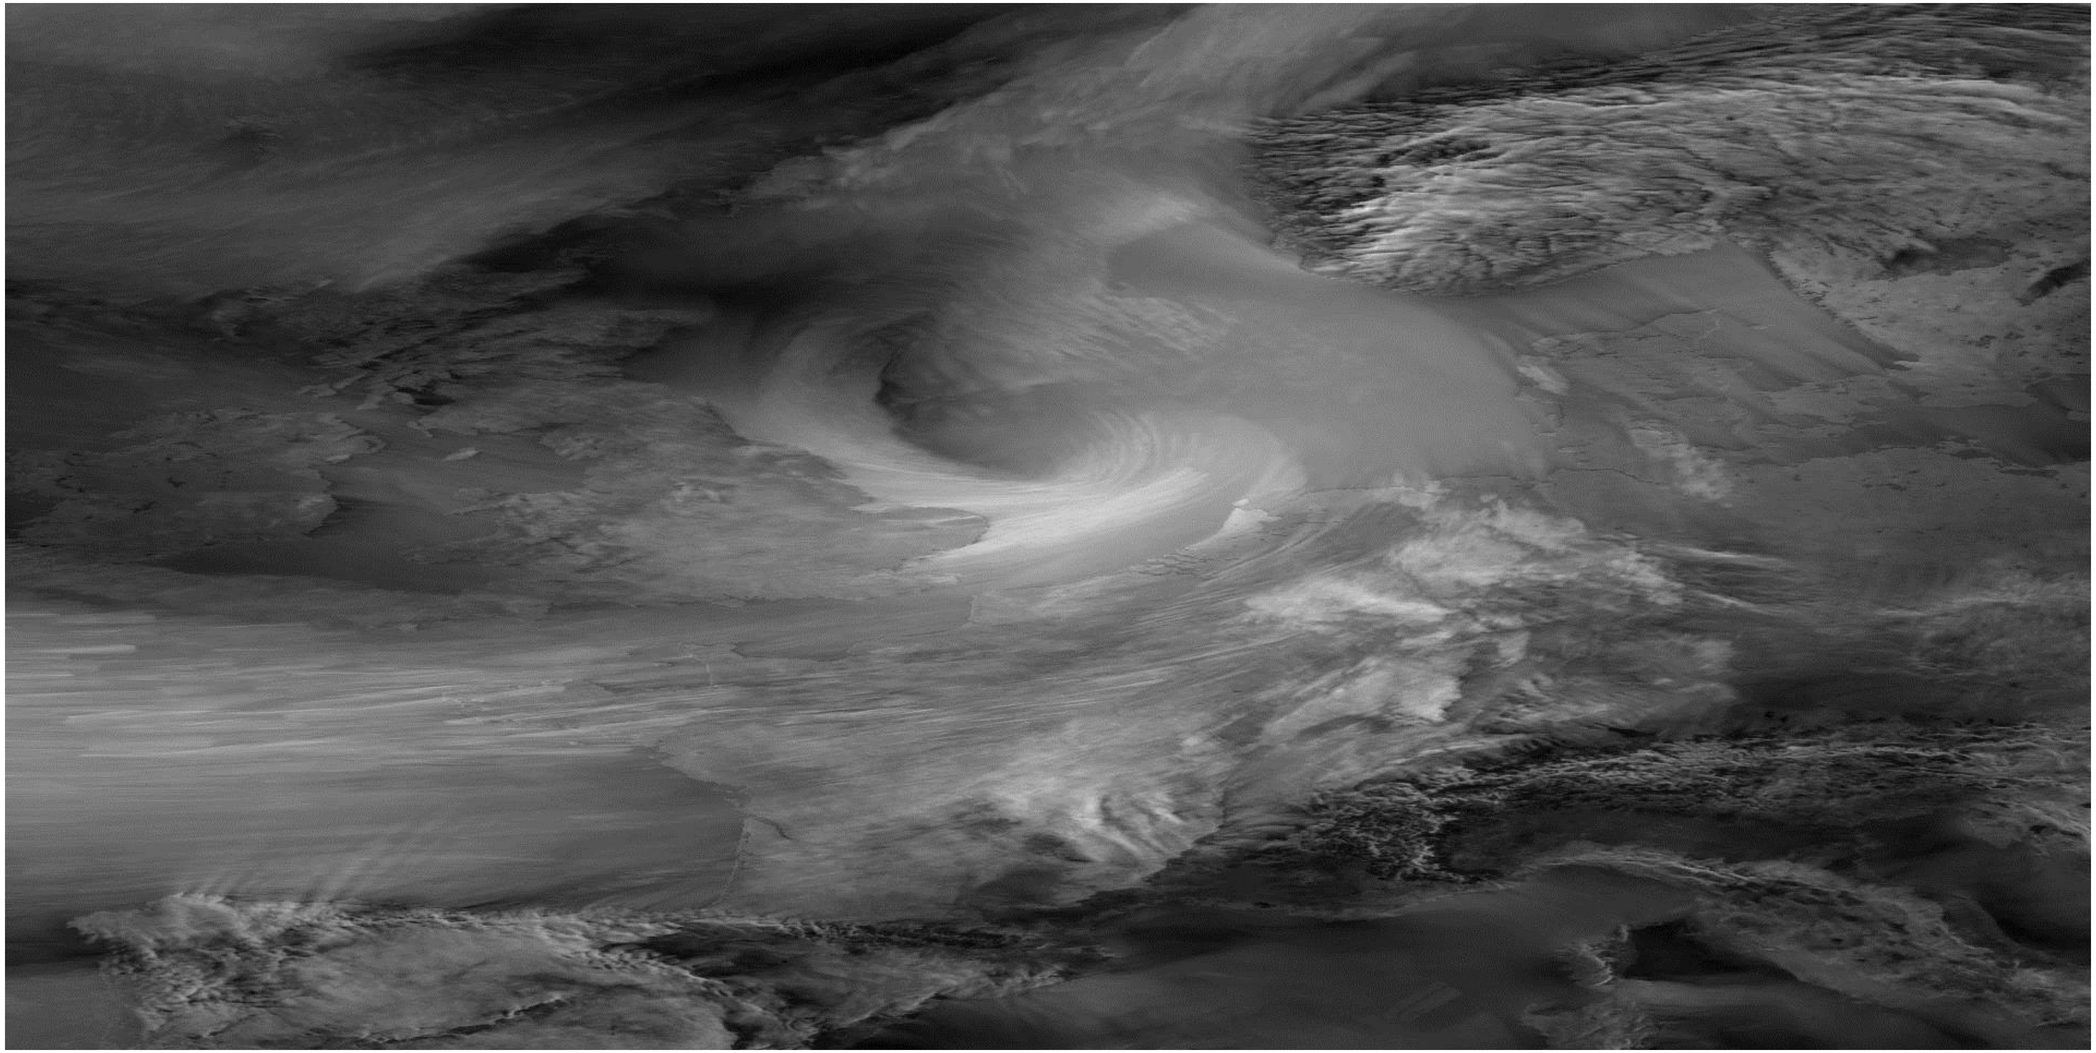

Conclusions

- High level of optimization in ECMWF’s IFS spectral transforms
  - 256 Megapixel camera with predictive skill!

... the global spectral transform method is not dead, but need to prepare to

- *Increase* the flexibility
  - discretization choices (and/or hybrid solution procedures)
- *Develop* alternative algorithms and methods
  - reduce data movement, communication and synchronization
- *Add* numerical and structural (code) flexibility
  - in the effort towards full Earth-System complexity

Additional slides

Local communication …

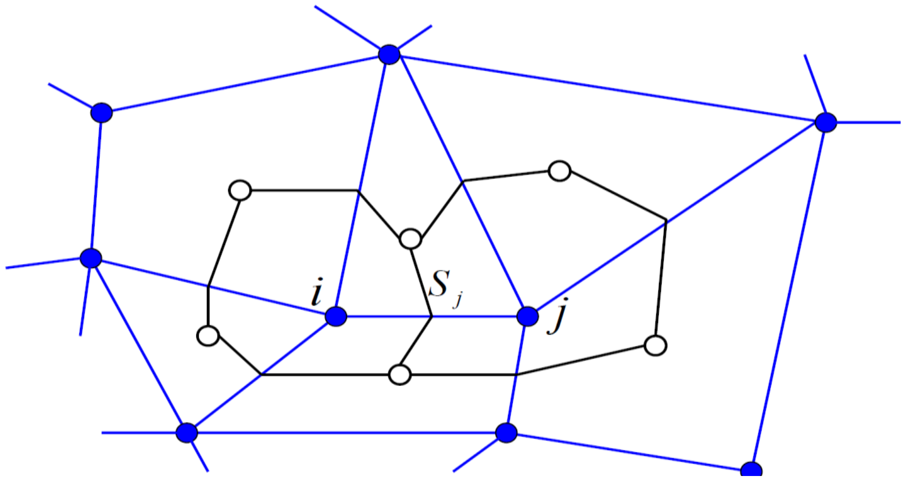

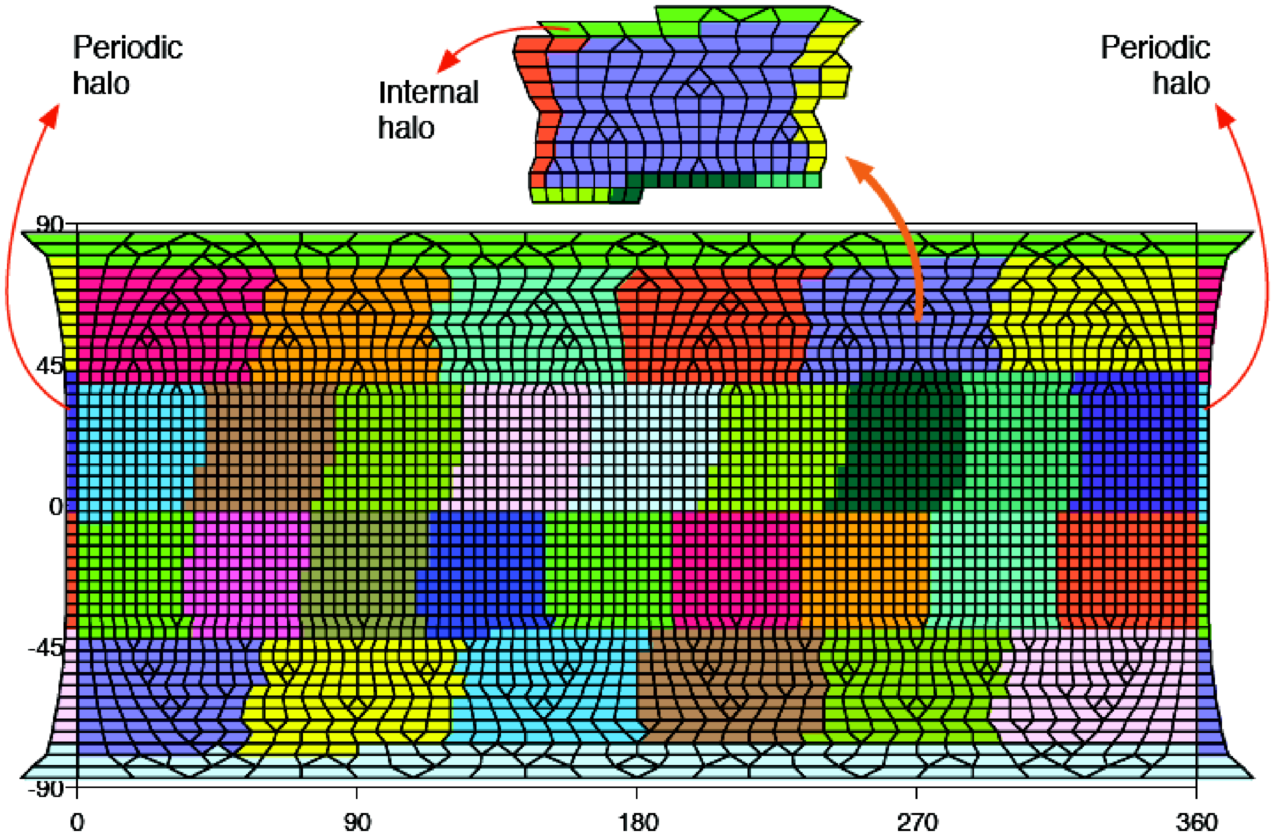

Day 15

Dry baroclinic instability, FVM (O640) versus the spectral IFS ($T_{co}639$):

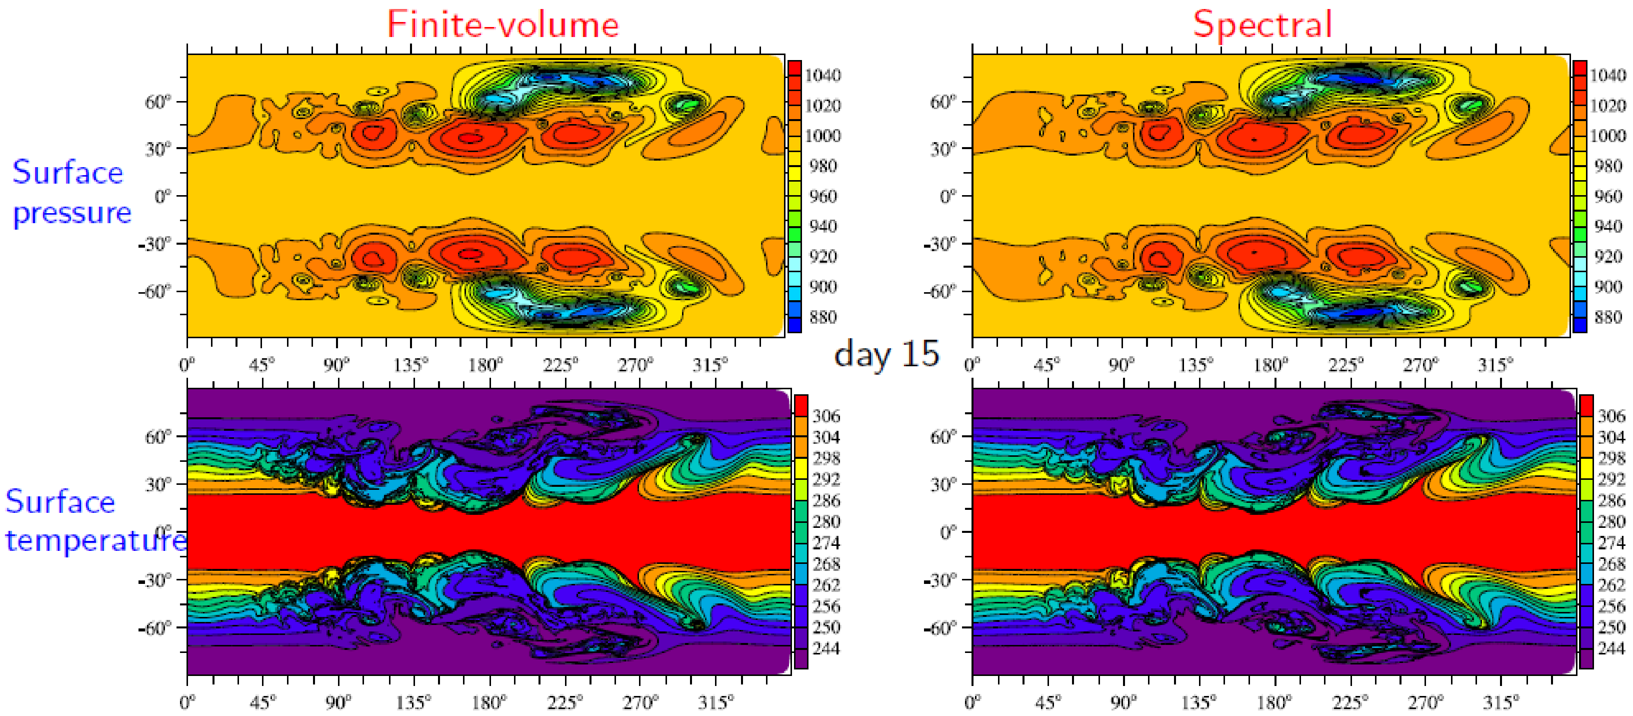

Application performance XC30

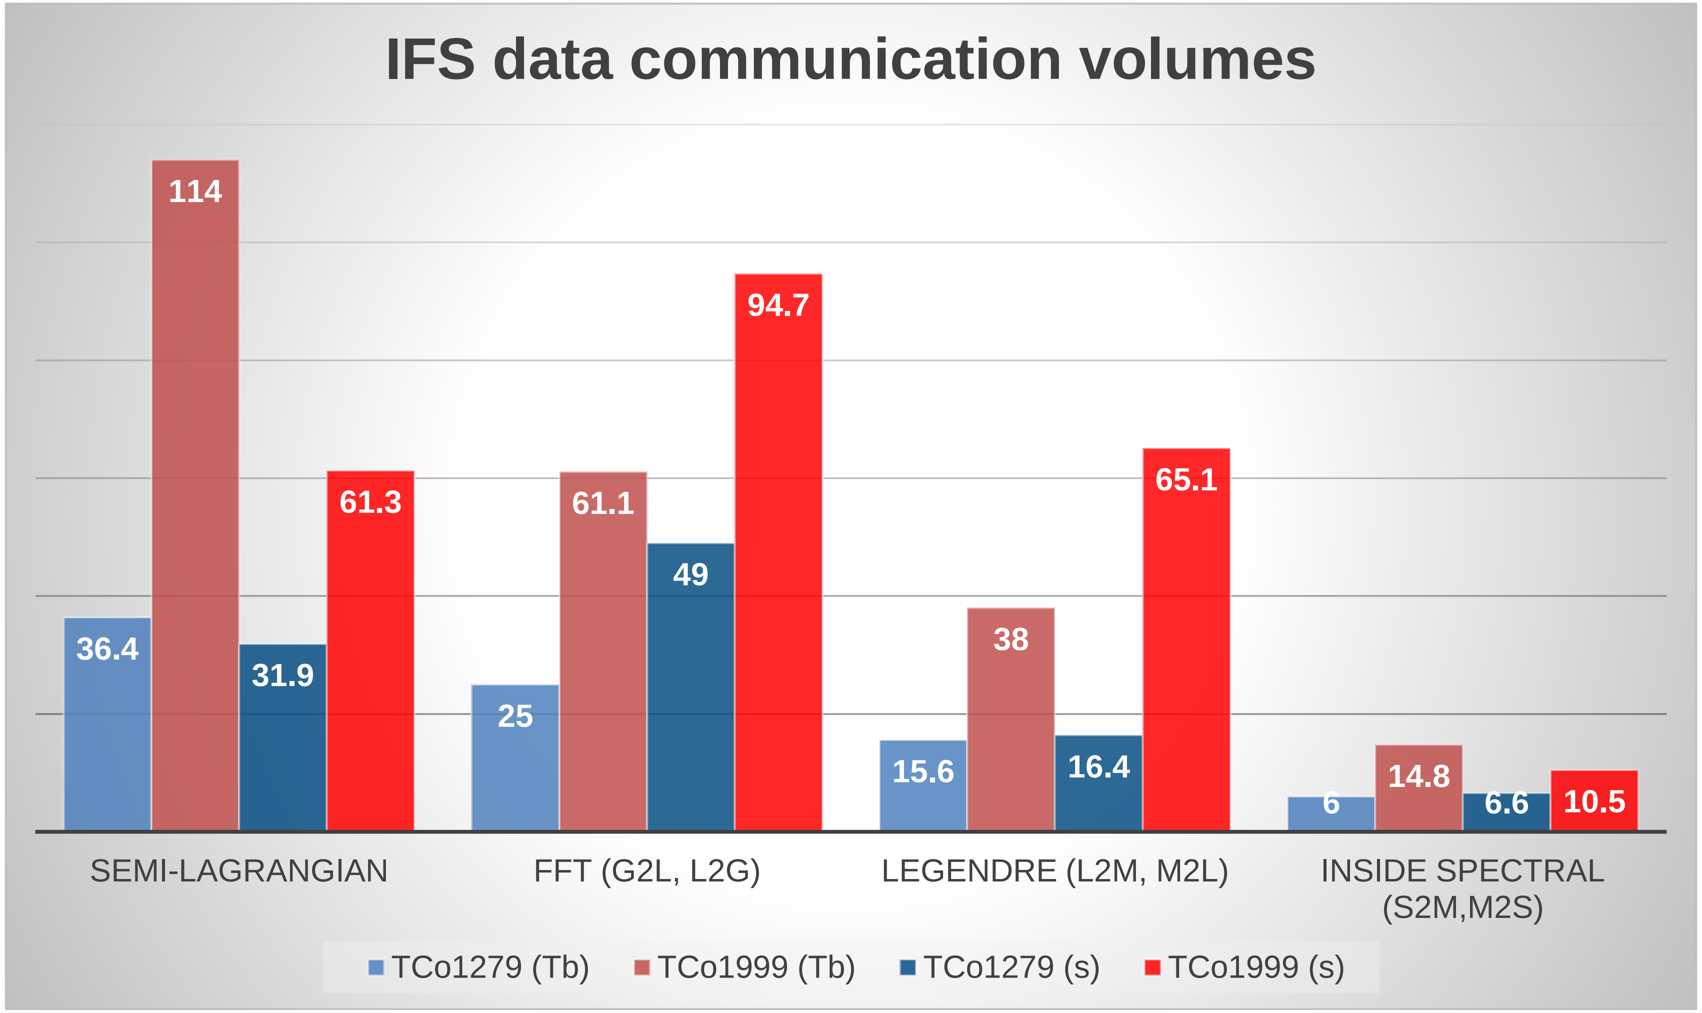

TCo1279 on 360 nodes; TCo1999 on 720 nodes ; dt=450s; 48h forecast

TCo7999 (~1.3km) total column liquid water (after 12h of simulation)

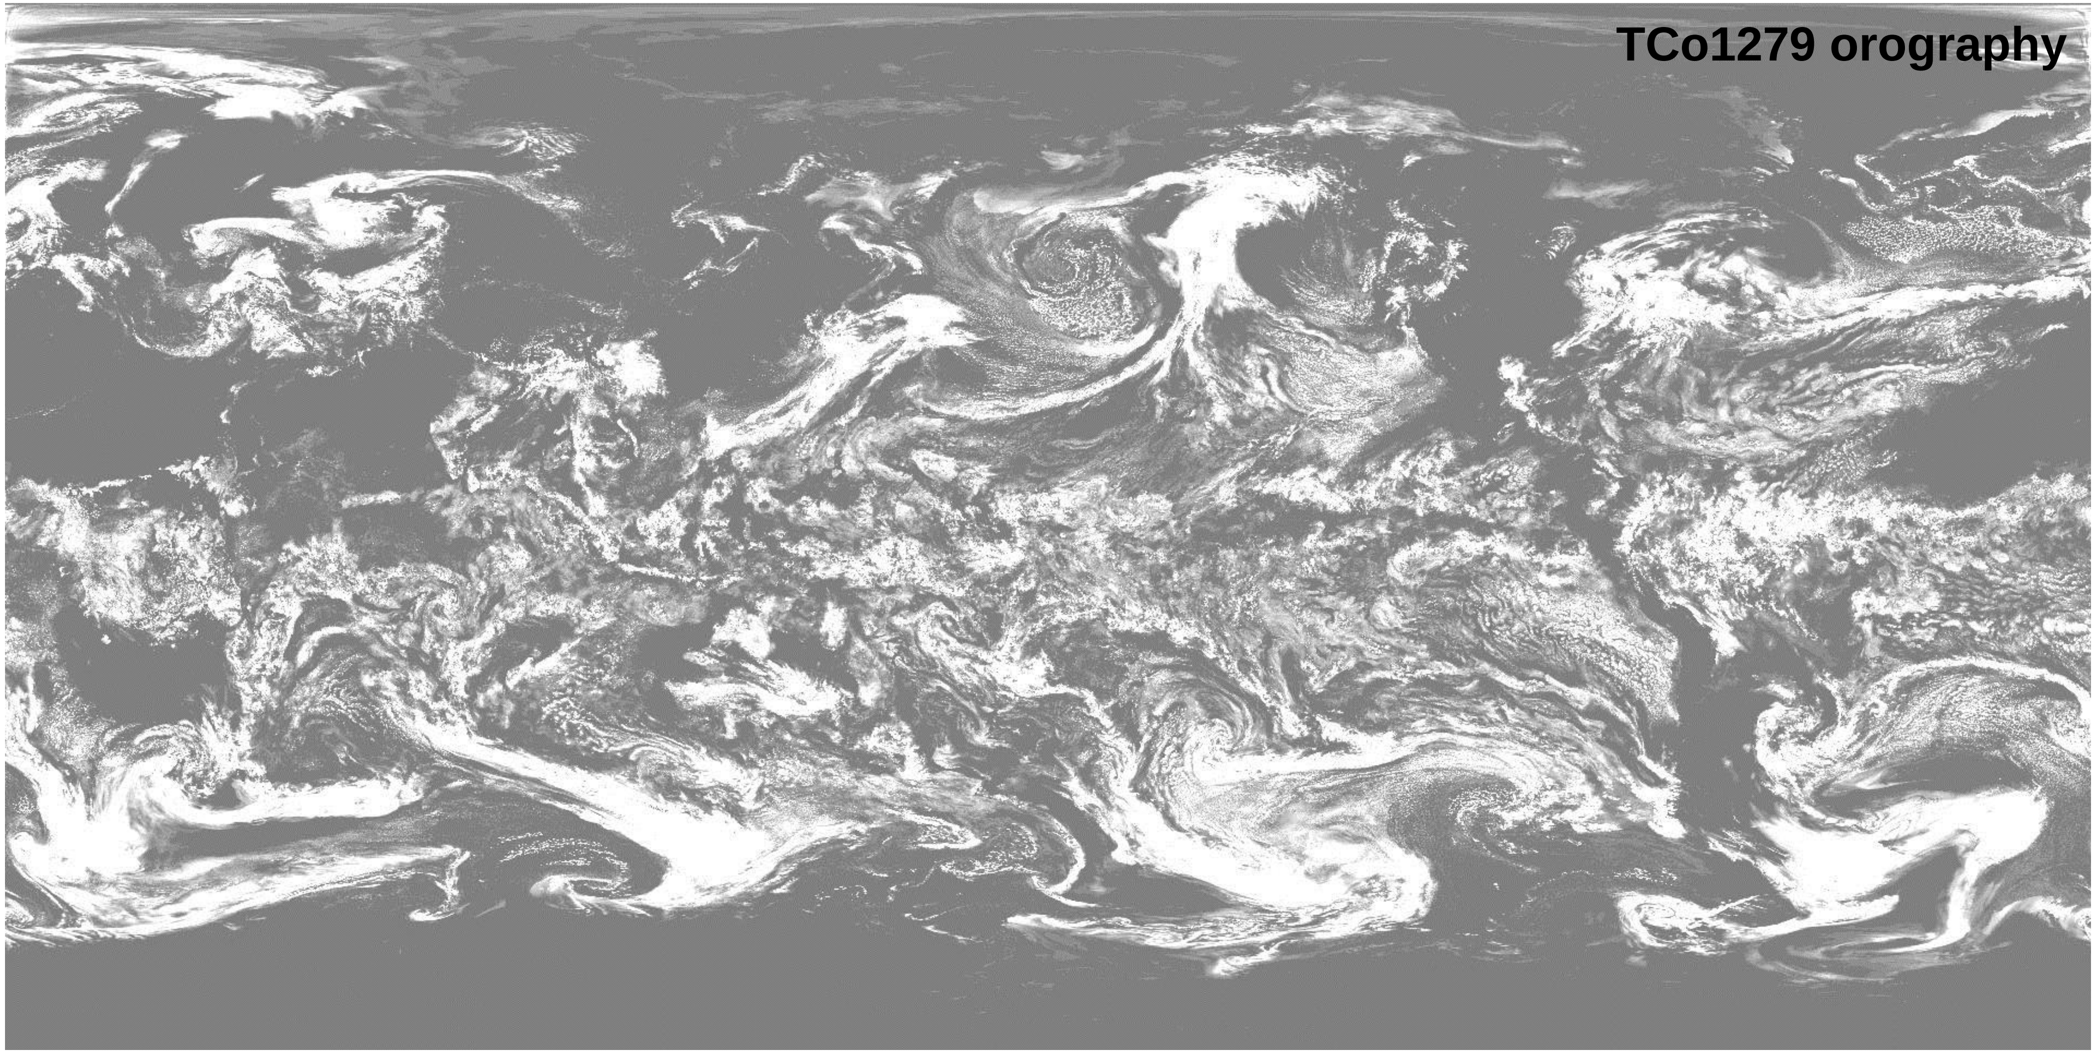

(12 h forecast, *non-hydrostatic*, no deep convection parametrization, 120s time-step, 960 Broadwell nodes, ~30s per timestep)