

ECMWF Feature article


.....
from Newsletter Number 146 – Winter 2015/16

METEOROLOGY

.....

A new grid for the IFS

.....



NASA Worldview

www.ecmwf.int/en/about/news-centre/media-resources

doi:10.21957/zwdu9u5i

This article appeared in the Meteorology section of ECMWF Newsletter No. 146 – Winter 2015/16, pp. 23-28.

A new grid for the IFS

Sylvie Malardel, Nils Wedi, Willem Deconinck, Michail Diamantakis,
Christian Kühnlein, George Mozdzyński, Mats Hamrud, Piotr Smolarkiewicz

ECMWF will implement a resolution upgrade for high-resolution forecasts (HRES) and ensemble forecasts (ENS) in spring 2016. HRES will then be run on a grid with a grid-point distance between neighbouring points of approximately 9 km instead of 16 km in the current configuration.

In the Integrated Forecasting System (IFS), many calculations are not carried out in grid-point space but in ‘spectral space’, where meteorological fields are represented by a sum of wave functions called spherical harmonics. As part of the resolution upgrade, ECMWF will move from a ‘linear’ to a ‘cubic’ grid by increasing the number of grid points used to represent each wavelength while keeping the number of spherical harmonics constant.

At the same time, it will use an octahedron-based method to reduce the number of grid points towards the poles. The resulting new ‘cubic-octahedral’ grid brings significant benefits in terms of computational efficiency and effective resolution.

From linear to cubic

The spectral transform method (Box A) has been applied successfully at ECMWF for more than 30 years, with the first spectral model introduced into operations at ECMWF in April 1983. The spectral transform method was introduced to numerical weather prediction (NWP) following the work of *Eliassen et al.* (1970) and *Orszag* (1970), who achieved high efficiency by partitioning the computations between a grid-point and a spectral representation at every time step.

In a spectral transform model such as the IFS, the horizontal wind, the temperature and the surface pressure have a discrete representation in two different spaces. In grid-point space, the representation of the fields is intuitive: each field is known by its value at each point of the model grid. The field discretisation in spectral space is more abstract. The idea is to represent the fields by a sum of analytical functions, the spherical harmonics, such that the sum of the spherical harmonics closely matches the field at each grid point. Each field is then known by the set of coefficients associated with the spherical harmonics in the sum (Box A). Each spherical harmonic has a characteristic horizontal wavelength, which is given by the value of its total wavenumber n . The wavenumber n indicates how many of the characteristic horizontal wavelengths are needed to go around the globe at the equator. The wavelength associated with the wavenumber n is given by $2\pi a/n$, where a is the Earth’s radius. The larger the wavenumber (the smaller the wavelength), the finer the scale represented by the spherical harmonic. The maximum wavenumber in the sum used to represent the meteorological fields, n_{MAX} , is the spectral truncation number of the model. The larger the truncation number, the smaller the scales potentially represented by the spectral approach. In spectral space, mathematical operations such as differentiation or integration are computed analytically using the series of spherical harmonics.

The accuracy of the transformation from grid-point space to spectral space and back is assured if the grid is a Gaussian grid, i.e. characterised by N specially determined quadrature points along a meridian between the pole and the equator, the ‘Gaussian latitudes’, and their associated ‘Gaussian quadrature weights’ used to compute the spectral coefficients (Box A). Several choices can be made to pair the maximum wavenumber of the spectral truncation, n_{MAX} , with the number of latitude circles between the pole and equator, N , which characterises the Gaussian grid.

In the current so-called linear grid, $n_{MAX} = 2N - 1$. A spectral transform using a linear grid represents the smallest wavelength $2\pi a/n_{MAX}$ by 2 grid points. Two other choices are the ‘quadratic’ grid and the ‘cubic’ grid, which represent the smallest wavelength by 3 and 4 points, respectively. The names ‘linear’, ‘quadratic’ and ‘cubic’ stem from the ability of the different grids to accurately represent linear, quadratic and cubic products in the equations.

The resolution of the IFS is indicated by specifying the spectral truncation n_{MAX} prefixed by the acronym TL (for triangular–linear), TQ (for triangular–quadratic), TC (for triangular–cubic) or TCo (for triangular–cubic–octahedral, see below). For example, the resolution of the current HRES is TL1279, a linear grid truncated at $n_{MAX} = 1279$.

Global spectral transforms

A

In the IFS, the horizontal wind, the virtual temperature and the surface pressure are transformed to spectral space and back to grid-point space at every time step. All the water variables and the passive tracers, e.g. specific humidity and prognostic cloud and precipitation, are kept in grid-point space because they are not needed for any of the computations in spectral space and because the spectral transforms can violate the positivity of the transformed field, in particular for non-smooth fields.

A direct spherical harmonics transformation of a field $\Phi(\lambda, \theta)$ known at each longitude λ_l and latitude θ_k is a Fourier transformation in longitude, followed by a Legendre transformation in latitude of the Fourier coefficients Φ_m at each zonal wavenumber m :

$$\Phi_m(\theta_k) = \frac{1}{L} \sum_{l=1}^L \Phi(\lambda_l, \theta_k) e^{-im\lambda_l} \quad \Rightarrow \quad \Phi_n^m = \sum_{k=1}^K w_k \Phi_m(\theta_k) \bar{P}_n^m(\theta_k)$$

where w_k is the Gaussian weight for the latitude k and \bar{P}_n^m are the normalised associated Legendre polynomials for the zonal wavenumber m and the total wavenumber n . In spectral space, the fields are known only by their set of spectral coefficients Φ_n^m .

The discrete Fourier transform is computed numerically very efficiently by using a Fast Fourier Transform (FFT). The discrete Legendre transforms

require the accurate discrete computation of the sum over the Legendre polynomials, which is accomplished by a Gaussian quadrature at the $K=2N$ special quadrature latitudes of the Gaussian grid between the two poles (the number of latitude circles between pole and equator, N , is used for the GRIB-encoded data).

The inverse discrete Legendre and Fourier transforms using a triangular spectral truncation n_{MAX} (i.e. $0 \leq n \leq n_{MAX}$ and $-n \leq m \leq n$) return the field in grid-point space, at each point of the Gaussian grid.

Recent concerns about the computational cost of the Legendre transform have been mitigated by a fast Legendre transform which exploits similarities of the associated Legendre polynomials to simplify

the computations (Wedi et al., 2013). Further computational acceleration can be expected from using modern hardware accelerator technologies. However, the parallel communications involved in the data transfer within transpositions from grid-point space to spectral space and back, at every time step of the model, remain a concern on future computing architectures (Wedi et al., 2015).

Higher effective resolution

Until 1998, a quadratic grid was used in the IFS to avoid the aliasing resulting from the computation of the Eulerian advection. With the implementation of the semi-Lagrangian advection scheme, a linear grid was introduced to enable finer scales in the spectral representation for a given grid resolution. However, recent experience suggests that the importance of non-linear processes increases with increasing resolution, thus exacerbating the problem of aliasing and requiring computationally expensive de-aliasing filters to suppress poorly resolved or misrepresented motions (Wedi, 2014).

The notion of resolution in an NWP model is more complex than just the grid spacing. In fact, selected processes are computed at coarser resolutions (e.g. radiation) and similar techniques may be applied in future to other physical processes. Moreover, the effective resolution of a model depends on the level of implicitly or explicitly applied filtering. Such filtering can be inherent in the equations used. For example, it would be impossible to resolve vertically propagating acoustic waves with a hydrostatic or an anelastic system of equations, whatever the time and space resolutions, as such waves are filtered by these equations. Some numerical algorithms may also be a source of significant filtering. Well-known examples are the effect of decentring in time-stepping schemes or damping in low-order interpolation. In addition, the filtering applied on boundary forcings such as the orography can be significant for the effective model resolution.

By definition, the cubic discretisation, where $n_{MAX} = N - 1$, filters wavenumbers between N and $2N$ (i.e. wavelengths between 2 and 4 times as long as the grid spacing). With such a discretisation, no numerical filter is needed apart from a small amount of numerical diffusion to eliminate any accumulation of energy at the wavenumber of the spectral truncation, and a damping layer at the model top to avoid artificial wave reflexion inside the computational domain. In addition, as shown in Figure 1, the cubic discretisation can stably support an orography with more variance in the small scales, thus providing the same spectral representation as that obtained from the original dataset for all wavenumbers almost up to the truncation limit. A cubic grid with a given truncation thus implies a higher effective resolution than a linear grid at the same and even at a higher truncation number.

As Figure 2 shows, the aliasing control required for the linear grid also has a significant impact on global mass conservation: the linear grid is associated with a steady increase in the mass conservation error over time while the cubic grid is nearly perfect without a mass fixer.

Consequently, the spring 2016 resolution upgrade (IFS cycle 41r2) will introduce a cubic spectral transform grid, with $n_{MAX} = 1279$ and $N = 1280$ latitude circles between the pole and the equator instead of $N = 640$ in the current system.

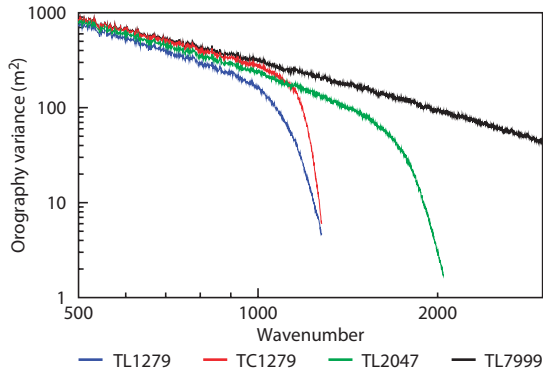


Figure 1 Variance of the orography as a function of the total wavenumber for different linear and cubic discretisations in the IFS.




Figure 2 Global mass conservation error for 10-day forecasts at different resolutions: TL1279 with a time step of $\Delta t = 600$ s, TC1279 with $\Delta t = 400$ s, TC01279 with $\Delta t = 400$ s. For comparison, TL2047 with $\Delta t = 400$ s is also shown. The chosen time steps are optimized for meteorological accuracy and efficiency at the given resolution.

Lower computational cost

Figure 3 shows the evolution of the computational cost of a forecast as a function of the grid resolution N . The blue ‘grid factor’ bars show how the cost would grow if it increased linearly with the number of points on the grid. It can be seen that the cost of a forecast using linear or cubic discretisations grows faster than the grid factor. However, for high resolutions the cost of the global communications in the spectral transforms significantly penalizes the linear grid compared to the cubic grid.

The greater computational efficiency of the cubic grid can also be seen in Figure 4, which shows the number of forecast days that can be produced per day depending on the number of available supercomputer cores. It can be seen that cubic discretisation requires a significantly lower number of cores to achieve the same number of forecast days than linear discretisation using the same physical grid.

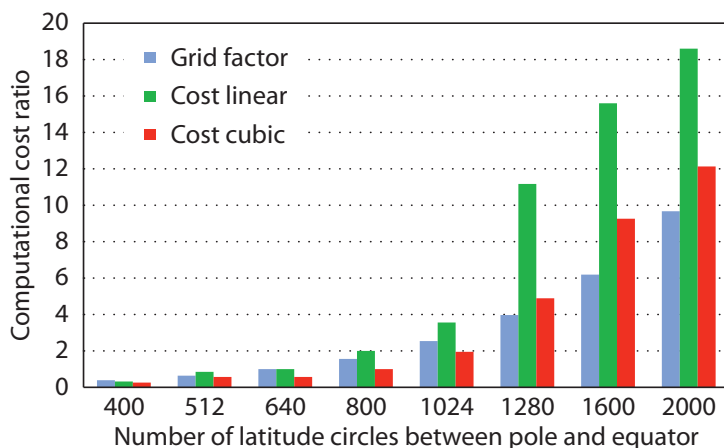


Figure 3 Computational cost relative to the current TL1279 as a function of the number of latitude circles between the pole and the equator, N . The green bars represent the relative cost of linear discretisations, the red bars that of cubic discretisations. The blue bars, shown for reference, represent the ‘grid factor’, the ratio between the number of points for a given reduced Gaussian grid divided by the about 2.1 million points at the current operational resolution TL1279 (i.e. a multiplication by about 4 of the cost for a doubling of N).




Figure 4 Computational capacity expressed in forecast days per day as a function of the number of supercomputer cores for the linear (TL2047) and the cubic (TC1023) discretisations associated with the same Gaussian grid $N = 1024$.

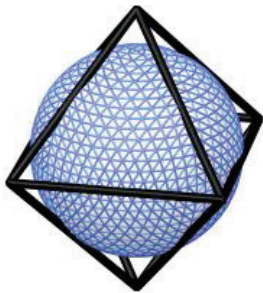
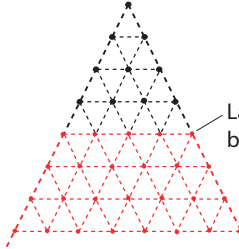
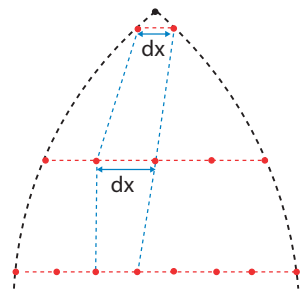
Octahedral reduced Gaussian grid

The choice of the spectral truncation and the number of latitude circles of the Gaussian grid is not sufficient to generate the IFS grid because it does not determine the number of grid points in the zonal direction at each latitude circle. If the same number of points is used for each latitude circle in a full Gaussian grid, the zonal resolution near the poles is substantially higher than the zonal resolution at the equator. Such a configuration generates a strong anisotropy of the discrete horizontal representation of the fields, potentially risks numerical instabilities and carries a significant computational cost due to the large number of points near the poles. For the past two decades, ECMWF has used a reduced grid (Hortal & Simmons, 1991), in which the number of points on each latitude circle is reduced towards the poles, keeping the relative grid-point distances approximately constant. This reduction lowers the number of points by approximately 30% without significant loss of meteorological accuracy in the spectral transforms.

A new method to reduce the number of grid points towards the poles has been explored for the next resolution upgrade of the IFS, both to optimize the total number of points around the globe and to introduce a regular reduction of the number of points per latitude circle towards the poles. The design of this new grid is inspired by a regular triangular mapping onto an octahedron, which corresponds to a reduction of 4 points per latitude circle, one per face of the octahedron (Box B). The resulting grid is called the 'cubic-octahedral reduced Gaussian grid'.

How to generate an octahedral reduced Gaussian grid **B**

1. Imagine each hemisphere of the globe is divided into 4 quarters, with each quarter corresponding to one face of an octahedron (left).
2. Start with 20 points, i.e. 5 per quarter, at the Gaussian latitude closest to the pole (middle).
3. Add one point per quarter for each new Gaussian latitude towards the equator, i.e. 4 more points per Gaussian latitude circle.
4. Because of the curvature of the Earth, the spacing between the grid points along a latitude circle varies with the latitude. It is slightly wider in the mid-latitudes than at the equator and near the pole (right).

The nominal resolution of the grid in the zonal direction is not as uniform around the globe as in the original reduced Gaussian grid (Figure 5), but the number of points per vertical level is significantly lower (for the $N = 1280$ grid, the original reduced grid has about 8.5 million points against 6.5 million for the new octahedral grid). In practice, this saves another 22% of total computation time. With the current reduced Gaussian grid, the number of points per latitude circle is constrained to be a multiple of 2, 3 and 5 by the Fast Fourier Transform (FFT) algorithm FFT992 originally developed at ECMWF by our former colleague Clive Temperton (Temperton, 1983). The cubic-octahedral reduced Gaussian grid will be used in the IFS together with the FFT package FFTW (<http://www.fftw.org/>), which efficiently allows any number of points per latitude circle.

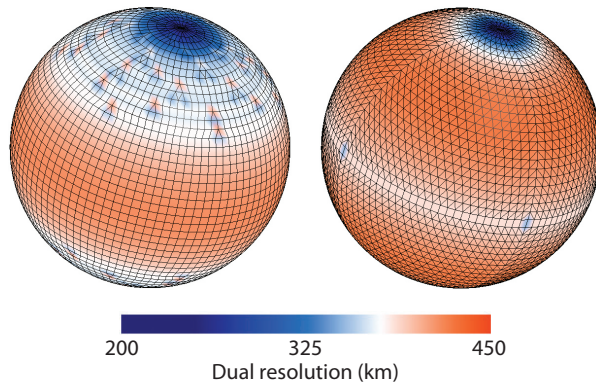


Figure 5 Primary meshes generated around the $N = 24$ reduced Gaussian grid points with 3.75 degrees (or approximate grid-point distances of 416 km). The shading represents the effective grid spacing of the median dual mesh, which is the mesh used for finite-volume computations, as explained by Smolarkiewicz et al. (2015). The shading is calculated as the square root of the local dual mesh area. The octahedral mesh (right) has a locally more uniform dual-mesh resolution than the original mesh (left). The coarse resolution is chosen for illustrative purpose only.

The spectral transform method may become computationally inefficient in the future due to the communication overhead of the global spectral transformations. To address this risk, alternative numerical methods that rely only on nearest-neighbour information are being developed at several Member State weather services and also in the context of the PantaRhei project at ECMWF (Smolarkiewicz et al., 2015). The recently investigated finite-volume module (FVM) can provide a more efficient and scalable way of computing differential operators in the IFS, but the accuracy of the approach depends also on the underlying mesh which defines the shape of the elementary volumes around which the computations are made. This is illustrated by results from numerical simulations of an idealised baroclinic instability using the FVM, as shown in Figure 6. Here, the more uniform mesh which is built around the grid points provided by the octahedral versus the original reduced Gaussian grid results in higher accuracy and substantially reduced unphysical flow distortions.

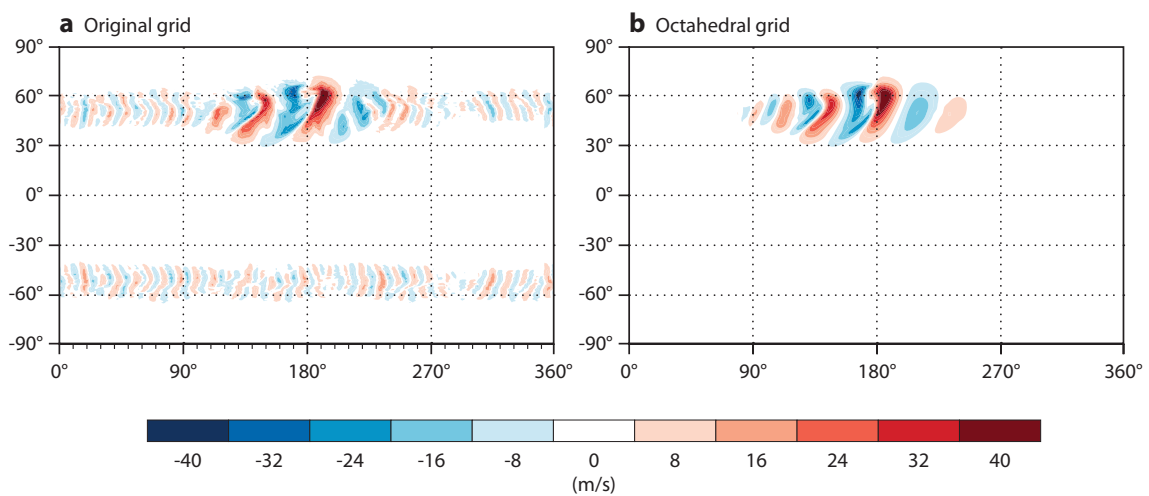


Figure 6 Idealised baroclinic wave test case using the finite-volume module of the IFS being developed at ECMWF, showing meridional wind component after 8 days of simulation for (a) the original reduced Gaussian grid and (b) the octahedral reduced Gaussian grid.

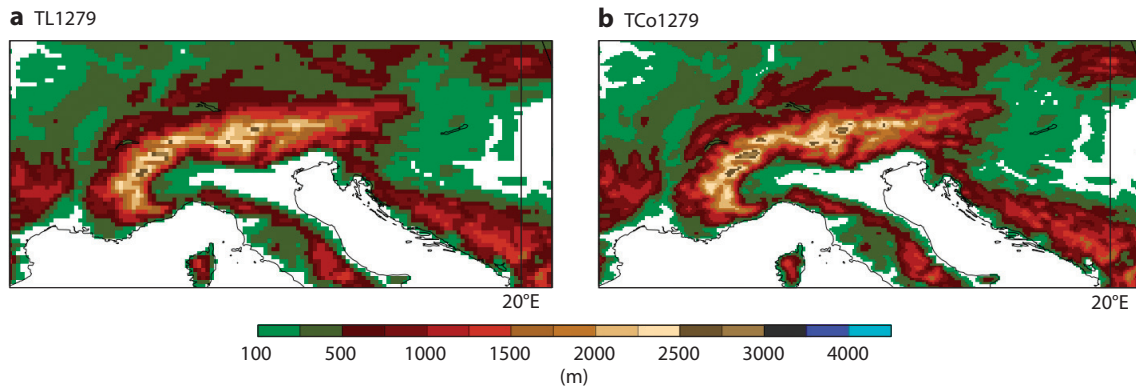


Figure 7 Representation of the Alps at (a) TL1279 and (b) TCo1279.

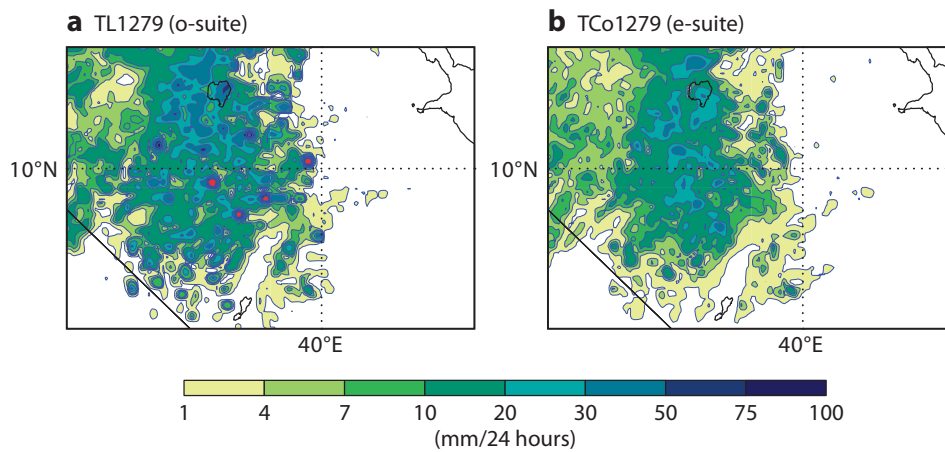


Figure 8 Accumulated 24-hour total precipitation for 24-hour forecasts starting from 15 July 2015 00 UTC in the Horn of Africa (Ethiopia, Djibouti) near Lake Tana for (a) the operational system (o-suite) at TL1279 and (b) the experimental system (e-suite) at TCo1279. The red dots in (a) show the location of spurious ‘grid-point storms’ where the predicted accumulated precipitation is larger than 100 mm.

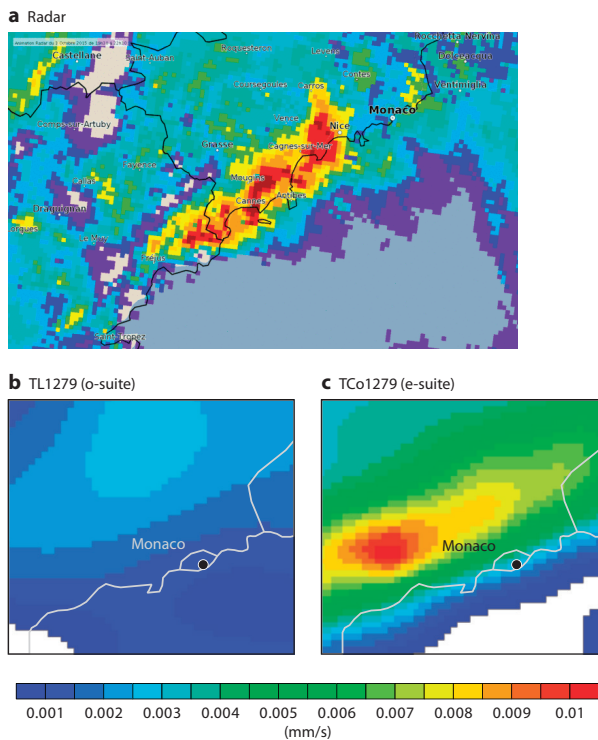


Figure 9 Total precipitation rate at 21 UTC during a severe event near Cannes on 3 October 2015 according to (a) Radar Mosaic observations, (b) a 3-hour forecast produced by the operational system (o-suite) at TL1279 and (c) a 3-hour forecast produced by the experimental system (e-suite) at TCo1279. The flash floods generated by intense and stationary convective systems during this event killed 20 people.

Summary of benefits

The cubic-octahedral spectral transform discretisation called TCo1279, with a spectral truncation at $n_{MAX} = 1279$ and a Gaussian grid with $N = 1280$ latitude circles between the pole and the equator, offers superior filtering properties at higher horizontal resolutions, an improved representation of orography, improved global mass conservation properties, and substantial efficiency gains. It also works well with more scalable, locally compact computations of derivatives and other properties that depend on nearest-neighbour information only.

With the cubic discretisation, the filtering of the spectral orography is reduced. As shown in Figure 7, the increased variance shown in Figure 1 directly translates into a sharper representation of the resolved orographic forcing in grid-point space.

Thanks to the optimized representation of smaller resolved weather features, problems such as ‘grid-point storms’ have been completely eliminated from the new system, as illustrated by Figure 8.

The resolution upgrade concerns not just HRES but also ensemble forecasts (ENS), the Ensemble of Data Assimilations (EDA) and the 4DVAR assimilation system. As a result of the cubic grid’s superior filtering properties, the level of filtering of the 4DVAR trajectory is now identical to that of the forecast model, and the initial analysis and evolving forecast have the same kinetic energy spectrum across the entire range of wavenumbers (not shown). Results of the evaluation being carried out in preparation for the resolution upgrade in 2016 indicate an improved effective resolution and exciting improvements to medium-range weather forecasts, as illustrated in a case study by Figure 9.

Details on the resolution upgrade can be found at
<https://software.ecmwf.int/wiki/display/FCST/Horizontal+resolution+increase>

Further reading

Eliassen, E., B. Machenhauer, & E. Rasmussen, 1970: On a numerical method for integration of the hydrodynamical equations with a spectral representation of the horizontal fields. *Report 2, Institut for Teoretisk Meteorologi, University of Copenhagen*.

Hortal, M. & A.J. Simmons, 1991: Use of reduced Gaussian grids in spectral models. *Mon. Wea. Rev.*, **119**, 1057–1074.

Orszag, S.A., 1970: Transform method for calculation of vector coupled sums: application to the spectral form of the vorticity equation. *J. Atmos. Sci.*, **27**, 890–895.

Smolarkiewicz, P., W. Deconinck, M. Hamrud, C. Kühnlein, G. Mozdzyński, J. Szmelter & N. Wedi, 2015: An all-scale, finite-volume module for the IFS. *ECMWF Newsletter No. 145*, 24–29.

Temperton, C., 1983: Self-sorting mixed-radix fast Fourier transforms. *J. Comput. Phys.*, **52**, 1–23.

Wedi, N.P., 2014: Increasing horizontal resolution in NWP and climate simulations – illusion or panacea? *Phil. Trans. R. Soc. A*, **372**, doi: 10.1098/rsta.2013.0289

Wedi, N.P., P. Bauer, W. Deconinck, M. Diamantakis, M. Hamrud, C. Kühnlein, S. Malardel, K. Mogensen, G. Mozdzyński & P.K. Smolarkiewicz, 2015: The modelling infrastructure of the Integrated Forecasting System: Recent advances and future challenges. ECMWF Research Department *Technical Memorandum No. 760*.

Wedi, N.P., M. Hamrud & G. Mozdzyński, 2013: A fast spherical harmonics transform for global NWP and climate models. *Mon. Wea. Rev.*, **141**, 3450–3461.

© Copyright 2016

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, England

The content of this Newsletter article is available for use under a Creative Commons Attribution-Non-Commercial-No-Derivatives-4.0-Unported Licence. See the terms at <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.