

The IFS cycle CY21r4 made operational in October 1999

On 12 October 1999 a substantial set of modifications to both the forecast model and the data assimilation system were introduced as the operational cycle (IFS cycle CY21r4). This article provides an overview over these changes, some of which are the result of several years of development. It is shown that the introduction of this new cycle has resulted in substantial improvements to the forecasting system. The changes that were introduced in CY21r4 are:

Data assimilation

- ◆ revisions to the background error statistics used in the 4D-Var system;
- ◆ revisions to the bias correction scheme for TOVS/ATOVS radiances;
- ◆ correction of the analysis of humidity from conventional data sources;
- ◆ the introduction of the assimilation of 10 m winds from the SSM/I instrument.

Model

- ◆ an increase in the number of model levels to 60, allowing higher vertical resolution, mainly in the planetary boundary layer;
- ◆ revisions to the cloud and convection schemes;
- ◆ improved fields for orography, land-sea mask, and the subgrid-scale orography scheme;
- ◆ revisions to the post-processing of 10 m wind.

These changes are discussed below and their major impacts highlighted. More detail on many of the modifications can be found in the papers listed at the end of this article and in ECMWF Research Department Technical Memoranda, which can be made available on request.

Data assimilation changes

New background error statistics

The technique used to derive the statistics of background (or first guess) errors is one of the most important aspects of any analysis system. To a large extent, these statistics determine the way in which information provided by the observations is interpolated and extrapolated to the model's grid points.

Since 1997, the ECMWF analysis has modelled the background error statistics by applying a sequence of steps that attempt to transform the model fields to a form in which the background errors have unit variance and are uncorrelated. The covariance matrix for errors in these transformed fields is simply the identity matrix, and the background cost function is consequently easy to evaluate. Each step of the transformation (often called the 'change of variable') requires the application of matrices or operators whose coefficients are determined empirically. The coefficients were derived using the so-called NMC method from a large sample of differences between pairs of forecasts of different duration verifying on the same date

(pairs of 48-hour and 24-hour forecasts were used). The sample of forecast differences provided a surrogate for a sample of background errors.

For CY21r4, the overall structure of the change of variable has been retained. However, the way in which the coefficients are calculated is different. An 'ensemble' of independent analysis experiments was conducted, each producing analyses for the period 3-28 February 1999 but which differed from one another because, for each experiment, the observations were randomly perturbed by an amount typical of the assumed observation errors. It can be shown that, if the model were perfect and the analysis system linear, and if the perturbations applied to the observations were statistically similar to the true observation errors, then the difference between any pair of background fields verifying at the same time would represent a sample from the distribution of background error. Differences between pairs of backgrounds produced in this way were used instead of forecast differences to calculate the background error statistics for the CY21r4 assimilation system. In order to remove the assumption of a perfect model, a representation of the effects of model error was included by perturbing the model during the first-guess forecasts using a 'stochastic physics' method.

The most obvious difference between the background statistics derived from the analysis ensemble and those calculated using the NMC method is that the former statistics have shorter horizontal and vertical correlation length scales. An illustration of this is given in Figure 1, which shows wave-number averaged vertical correlation matrices for vorticity background errors calculated by the two methods. The NMC method gives much deeper vertical correlations, particularly in the middle-to-upper troposphere.

Revised bias correction of TOVS/ATOVS radiances

In the past, the RTOVS bias correction files were recalculated every month. A new method, using model predictors, was introduced in IFS-cycle CY18r6 and, because it became clear that the bias correction produced in this way was more stable, the updating period was extended to about two months. When ATOVS 1C radiances replaced RTOVS the IFS-cycle CY21r1 used a bias correction computed 4 months earlier. Prior to the introduction of CY21r4, an extensive review was undertaken of the bias correction procedures, taking into account the new stratospheric levels that had been introduced in the earlier 50-level model and the availability of data from the new AMSU instrument. As a result it was decided to use six model predictors and to revise the geographical sample-matching and statistical screening. The new procedure was tested with the CY21r4 model and was shown to improve the forecast performance of the system.




Figure 1: Wave-number averaged vertical correlation matrix for vorticity background errors calculated using (a) the new method and (b) the NMC method. Model levels 30, 39 and 49 are near pressure levels 200 hPa, 500 hPa and 850 hPa, respectively. The vertical correlation length scales are shorter in (a), especially in the middle and upper troposphere.

Correction of the humidity analysis

SYNOP and radiosonde humidity observations are received in the form of dew-point temperature data. For ease of use in the data assimilation, these data are converted to either relative humidity (SYNOP) or specific humidity (sondes). However, the way that this conversion has been done for many years at ECMWF has not followed WMO guidelines. The conversion ought to be done using the saturation vapour pressure calculated with respect to water at all levels in the atmosphere, without involvement of ice or mixed ice/water phases. However, prior to CY21R4, the conversion had been done using the model's definition of saturation vapour pressure, which allows for ice and mixed phases. The use of the correct procedure leads to an apparent increase in measured humidity, especially where the relative humidity is high and the temperature is far below freezing. In areas with good radiosonde data-coverage, the effect on data assimilation is a substantial moistening of the upper troposphere. In zonal-mean terms the CY21r4 analyses are between 6% and 12% moister than those produced by the previous scheme, especially in the 500-300 hPa layer north of 40°N, as shown in Figure 2. In the tropics, and in the Southern Hemisphere, the area-averaged impact is much smaller due to the lesser density of the radiosonde network.

With the corrected treatment of humidity observations there is a better agreement between observed and model humidity, and this results in reduced humidity analysis increments in the assimilation. Previously the observations appeared to be biased dry in the upper troposphere compared with the model, whereas now they appear to be largely unbiased. The increased moisture has also affected the amount of high cloud in the model, with an increase from around 27% to 32% cloudiness in the Northern Hemisphere mid and high latitudes. The forecast impact was tested over the 40-day period from 26 July to 4 September 1999, and there was a surprisingly large and positive improvement for the Northern Hemisphere troposphere.

Over a timescale of several weeks the test assimilation increased moisture in the lower stratosphere in high and mid latitudes. In subsequent operational use, the lower-

stratospheric humidity became quite unrealistically high, leading to significant systematic forecast errors in temperature and eddy kinetic energy. A cure for this side effect of the corrected treatment of humidity observations has been implemented in IFS cycle CY22r1. The operational analyses produced from 12 October 1999 to April 2000 with CY21r4 must, however, be regarded as being unsuitable for use in studies of local stratospheric humidity.

The post-processing of the 2 m dewpoint has also been changed to reflect WMO guidelines for observing and reporting, i.e. to use saturation vapour pressure with respect to water only (N.B. users of the analysed and forecast 2 m temperature and dewpoint to infer near-surface relative or specific humidity will need to change their procedures accordingly).

10 m winds from SSM/I

Estimates of total column water vapour derived by a one-dimensional variational analysis (1D-Var) of SSM/I radiance data have been assimilated operationally at ECMWF since 29 June 1998. The 1D-Var also produces estimates of surface wind speed over the sea; these have been monitored for some time and have been found to be of good quality. Consequently, experiments have been conducted to test the impact of assimilating these SSM/I wind speed estimates in the ECMWF 4D-Var analysis system. Using data from the DMSP-F13 spacecraft, it has been found that the assimilation of the SSM/I wind speed estimates generally increases the model surface wind. The largest changes are found over the southern oceans in winter where the SSM/I data cause mean wind speed increases of up to 0.5 m/s over large areas. Smaller, but still significant, wind speed increases are found over the North Atlantic and North Pacific in winter. The systematic changes to the analysis have resulted in a better agreement with conventional marine wind observations (from ships and buoys), but have also improved the analysis fit to independent (i.e. not assimilated) ERS-2 altimeter data (Figure 3). Furthermore, a number of cases have arisen when the use of the SSM/I wind speed data has had a major synoptic impact upon the analysis and has produced significantly improved medium-range

forecasts. As a result of these tests, the SSM/I wind speed estimates are now assimilated operationally in the CY21r4 system.

Model changes

Vertical resolution

The impact of the increased stratospheric vertical resolution introduced with the earlier 50-level version of the model has been described in a Newsletter article by Untch *et al.* (ECMWF Newsletter No. 82, Winter 98/99, pp. 2–8). The increased number of levels to 60 in the CY21r4 model doubles the vertical resolution below 1500 m; adequate vertical resolution in the lower troposphere has been shown to be fundamental for the proper representation and prediction of boundary-layer processes. This brings the lowest model level down to 10 m above the surface compared with 33 m in the previous 50-level model. Other consequences of the new vertical resolution are that the level of poorest vertical resolution moves up from about 730 hPa to 575 hPa, and there are three additional model

levels above 850 hPa in the troposphere and one additional level in the stratosphere.

A schematic of the distribution of levels for the L50 and L60 versions from the surface to 700 hPa (where the largest increase in resolution is located) is shown in Figure 4. In long integrations (120 days), the low-cloud cover in the L60 configuration has generally increased relative to the L50 model, particularly over the stratocumulus areas and the southern ocean – an improvement according to ISCCP observations. The L60 model also produces more realistic precipitation fields in the oceanic equatorial regions.

L60 and L50 assimilation experiments have also been performed for two periods in winter and summer. The L60 model fits the TEMP humidity observations better than the L50 model globally, and a positive impact from the L60 version is also clear from the statistics of the tropical TEMP temperature observation increments in the troposphere.

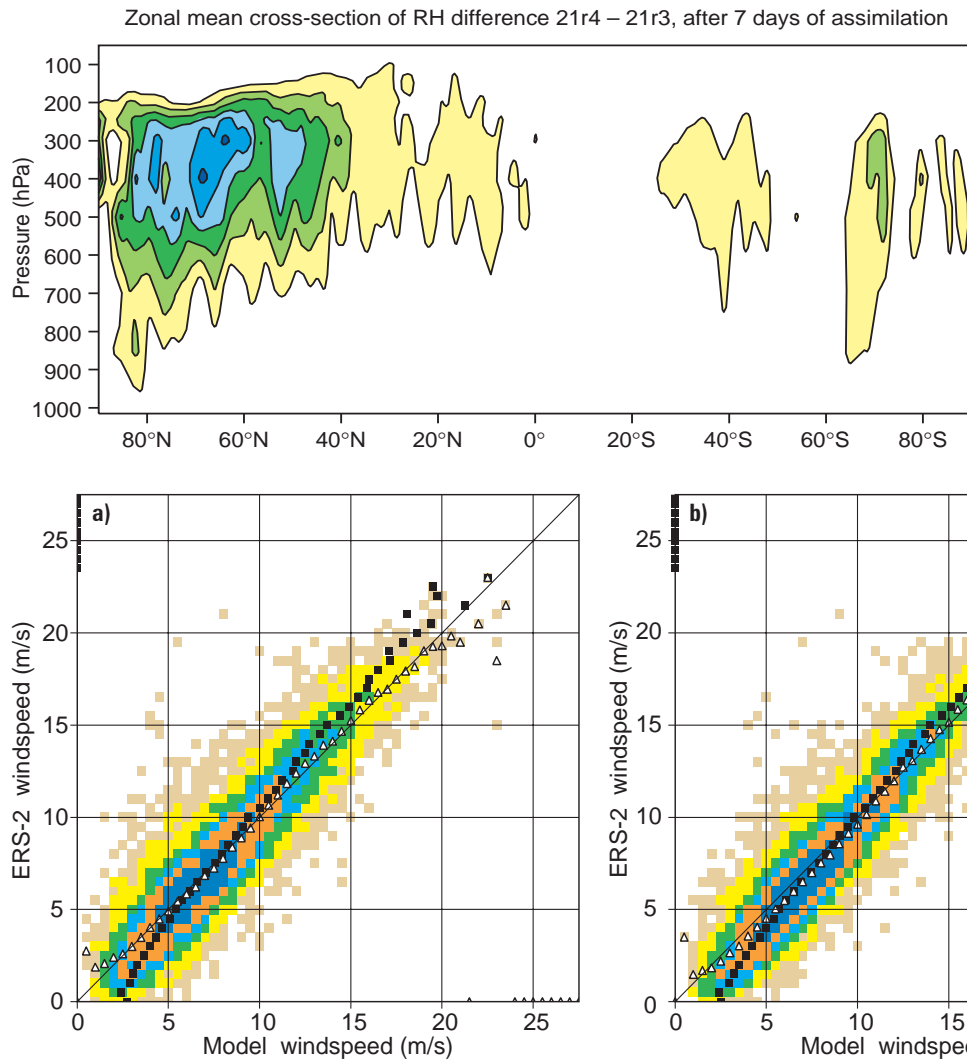


Figure 2: Relative-humidity differences after 7 days of assimilation between systems with and without the correct procedure for converting radiosonde and surface humidity observations. Note the apparent increase in humidity in the upper troposphere in regions of relatively dense radiosonde observations.

Figure 3: Fit of analysed 10 m wind to ERS-2 altimeter data for analyses from 1500 UTC 15 September 1998 to 0900 UTC 30 September 1998 when (a) the SSM/I wind speed estimates are assimilated, and (b) when they are not. The use of the SSM/I winds in the analysis reduces the wind bias from -0.35 m/s to -0.03 m/s.




Figure 4: The distribution of model levels below 700 hPa for the 50-level (left) and 60-level (right) versions of the ECMWF model.

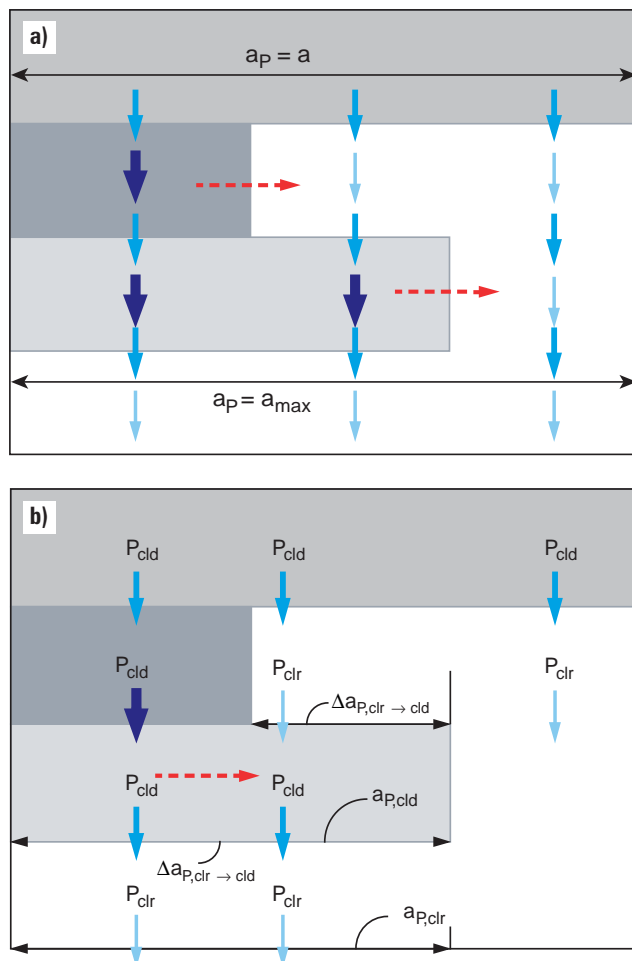


Figure 5: Schematic of the treatment of cloud and precipitation overlap in (a) the old cloud parametrization and (b) the new cloud parametrization. The vertical arrows denote falling precipitation, with P_{cld} and P_{clr} signifying the precipitation in the cloudy and clear regions, respectively. The horizontal dashed arrows indicate the apparent transfer of water arising as a consequence of the averaging process.

Cloud and convection scheme changes

A large number of changes to both the cloud and the convection parametrization have been introduced into the CY21r4 operational model. At the heart of the changes to the cloud scheme is a revision of the treatment of precipitation. In previous versions of the ECMWF model, as in most other GCMs, precipitation was described by using a grid-mean flux, with a possible separation into rain and snow. Recent research has identified serious shortcomings in this approach if the fractional area covered by clouds within a grid box varies with height. The main reason for this is illustrated in Figure 5(a), which shows four model layers, three of which are covered by clouds of various sizes (grey boxes). Precipitation is generated in the top cloud level. As it falls into the layer below, part of it enters cloud and part falls into clear sky. The part inside the cloud will be enhanced due to conversion and/or collection processes whereas, outside the cloud, precipitation will evaporate. If only the grid-mean flux is used to describe precipitation, the two very different precipitation fluxes (large inside cloud, small outside cloud) need to be averaged before entering the next layer. This implies a horizontal water transport from cloud into clear sky (indicated by a dashed horizontal arrow). Due to this transport, more precipitation is available for evaporation in the clear-sky part of the next model layer. This has been shown to lead to a substantial overestimation of evaporation in the presence of vertically varying cloud fraction. In order to address the problems outlined here, a new treatment of precipitation has been developed that separates precipitation fluxes into cloudy and clear-sky contributions, in a very similar way to the treatment of radiative fluxes in cloudy columns. The basic principle of this scheme is illustrated in Figure 5(b). The separation of cloudy and clear-sky precipitation avoids the averaging problems between the two parts of the grid box, although averaging within each of the two portions of the grid is still necessary (as indicated in the cloudy part of the third layer from the top in Figure 5(b)). However, the effects of this averaging are small compared with averaging between cloud and clear sky. An interesting “by-product” of this new treatment of precipitation is a prediction of what fraction of the grid box is covered by precipitation (often referred to as the precipitation fraction). The availability of this product opens up new possibilities for precipitation forecasting and validation, and these are currently under investigation.

For brevity, all other changes to the cloud parametrization are only listed here. They consist of:

- ◆ a revised treatment of the cloud sources due to convection;
- ◆ a revised treatment of stratiform cloud generation;
- ◆ a separation of small and large ice particles in the description of ice sedimentation;
- ◆ the introduction of an implicit numerical scheme for precipitation evaporation;
- ◆ a change to the treatment of threshold relative humidity above which no precipitation evaporation takes place;

- ◆ an increase in cloud erosion in shallow cumulus clouds;
- ◆ a revised treatment of the pure-ice mixed-phase transition layer.

Various modifications to the convection parametrization have also been introduced into the new model. They are:

- ◆ an increase of the relaxation timescale in the closure for deep convection to a minimum value of one hour at horizontal resolutions higher than TL159;
- ◆ a reduction in the strength of the penetration of the tropical tropopause by deep convective updraughts;
- ◆ an enhancement of the conversion from liquid water to precipitation by a factor of 1.3.

New orography fields

Global model orography, land-sea mask, and subgrid-scale orography fields have been created for each model resolution from a more detailed orography and land-cover dataset. The so-called USNAVY dataset containing terrain height and percentage of land at 10' x 10' resolution (about 15 km) was used to generate these fields which were included in the operational model from 1 April 1981 to 1 April 1998. Subsequently, for CY18r5 and later cycles, the global orography and land-sea mask (but not the subgrid orography fields) were calculated from a new 2'30" x 2'30" dataset created by Météo-France in co-operation with ECMWF. In CY21r4, a further change has been made to use a newer dataset for terrain heights at a finer resolution of 30" x 30" (GTOPO30), together with a dataset for land cover with similar resolution, i.e. about 1 km global resolution. The GTOPO30 data, distributed by the US Geological Survey EROS Data Centre ⁽¹⁾, has been combined with a special dataset for Greenland. The land-cover dataset is based on the processing of two years of 1 km AVHRR data, and is distributed by the US Geological Survey EROS Data Centre ⁽²⁾. The spectrally unfitted orography is obtained by an aggregation of the GTOPO30 data on the model grid. Since the GTOPO30 data comes with an ocean mask (lakes are labelled as land) the land-sea mask is obtained from the 'water' type land-cover dataset. The difference between the 'water' mask and the 'ocean' mask is used to build a 'fresh-water', or 'lake', mask that has been employed as an auxiliary field in the SST analysis for CY21r1 and later cycles.

Apart from being based on the different reference orography data, the computation of the fields in CY21r4 differs from the previous algorithm in a number of ways. The subgrid-scale orography fields are calculated such that the scales below about 5 km do not contribute (these smallest scales will contribute to new roughness-length calculations that are planned for the future). Also, the derivatives needed for computing some of the subgrid fields are equal-area derivatives, and the slope of the resolved model orography is removed from the data before the calculations are applied.

Since there are substantial differences from the previous subgrid-scale fields, especially for the slope (which is sometimes larger by up to a factor of three), it was crucial to test thoroughly their impact on the model and to assess the need to adjust the parameters in the associated parametrization scheme. Overall, it was found that the impact of merely replacing the fields was positive in terms of systematic errors in seasonal ensembles of T63 experiments, and so it was decided not to try to re-tune the parametrization scheme. The results of tests using the new orography fields in medium-range forecasts showed a positive impact on forecast scores for the Northern Hemisphere and neutral impact for the Southern Hemisphere.

Finally, a further change has been made to correct two small coding errors that were identified as a result of the adaptation of the ECMWF subgrid orography scheme to the new DWD model. No significant impact of these corrections was found either in extended integrations or in medium-range forecasts.

Post-processing of 10 m winds

The observing stations used for the verification of weather parameters are usually located in open areas, whereas the model roughness length is a so-called 'effective roughness' chosen to obtain adequate area-averaged momentum fluxes. However, the momentum flux (and, therefore, also the 'effective roughness' length) is dominated by the rough terrain features (due, for example, to orography and patchy high vegetation). Consequently, the resulting area-averaged 10-m model wind tends to be lower than observed in open sub-areas. In order to convert the model winds to a form that is comparable to typical 10-m wind observations an 'exposure correction' is applied during the post-processing. Prior to CY21r4, the model winds were interpolated (over land only) from the lowest model level (at about 30 m in the L50 model) to a height of 10 m using a local roughness length of 0.03 m, if this is lower than the model roughness length. The idea was that there is a height above the surface (the blending height - set to the height of the lowest model level, for convenience), where the effects of surface heterogeneity have merged. Below the blending height, the wind profile adapts to the local terrain. In the L60 model, with the lowest model level at 10 m, the mechanism of exposure correction is no longer appropriate. Consequently, the blending height has been set to a value of 75 m, independent of the model's resolution. The wind speed is interpolated to this level from the model levels, followed by a further interpolation to the 10-m level using the same local roughness-length assumption as before. However, the wind direction at 10 m is taken to be the same as that of the model wind at the lowest level. Comparisons between 10 m wind speeds forecast by the L50 and L60 models indicate that the results are improved in the L60 model, the wind speeds being, in general, slightly higher and closer to the observations.

(1) <http://edcwww.cr.usgs.gov/landdaac/gtopo30/gtopo30.html>

(2) <http://edcwww.cr.usgs.gov/landdaac/glcc/glcc.html>

The main impacts of the model changes

Cloud cover

The impact of the changes on the model’s cloud fields is summarised in Figure 6. This shows the zonal mean of the total, low-level, mid-level and high-level cloud cover taken from 5-day forecasts for the entire month of September 1999 from the L50 model and the CY21R4 system (which was running experimentally at that time). The most obvious change is a marked increase in low-level cloud cover by about 10% occurring at all latitudes. The main cause for this change has been found to be the interaction of the cloud and convection parametrization with the increased boundary-layer resolution. The high-level cloud cover is reduced everywhere, particularly in tropical areas. There is also a slight reduction in mid-level cloud cover. Recent comparisons of the model cloud cover with retrievals using data from the HIRS instrument on board the polar-orbiting satellites indicates that the changes in high-level and low-level cloud cover are improvements. The net effect on total cloud cover is a slight reduction in the tropics and a small increase in the extratropics. The change in vertical structure is, however, much more significant.

Low-level relative humidity

The changes in the treatment of the evaporation of precipitation lead to substantial changes in relative humidity. There is a decrease in the zonal-mean relative humidity of 3% to 5% in the tropical mid-troposphere. The near-surface relative humidity is generally increased by about the same amount. Table 1 shows the background departures of relative humidity at the 2 m height compared

with SYNOP observations, as measured in the data assimilation process. Both the mean and the root mean square of the difference between the observations and the model’s first guess are shown, averaged over stations in the Northern and Southern Hemisphere and in the tropical belt. It is evident that a dry model bias (positive background departure) has been largely alleviated in all regions together with a reduction in the RMS difference. This indicates that the new model version predicts more realistic values of low-level relative humidity. The higher values of low-level relative humidity lead to a reduction of a positive bias in surface evaporation over much of the tropical and subtropical oceans.

		Summer		Winter	
		Bias	RMS	Bias	RMS
N. Hemisphere	Old	4.2	11.1	1.3	10.6
	CY21r4	2.1	10.3	-0.2	10.6
Tropics	Old	2.4	12.0	1.9	10.8
	CY21r4	0.3	11.0	-0.2	9.7
S. Hemisphere	Old	3.4	13.0	2.4	11.3
	CY21r4	0.8	11.5	1.7	10.2

Table 1: Bias and root mean square (RMS) differences from the background field (observations minus first guess) for 2 m relative humidity (%) for the previous operational system (Old) and the CY21r4 system. The summer results are from six data assimilation cycles in June 1999; the winter results are taken from 16 cycles in January 1999.




Figure 6: Zonal mean distribution of (a) the total, (b) the low-level, (c) the mid-level, and (d) the high-level cloud cover taken from all 5-day forecasts for September 1999 with the then operational system (red line – OPS) and the new CY21r4 system (black line – NEW).

Tropical temperatures

Another significant impact of the new model is a change in the vertical temperature structure in the tropics. This is summarised in Figure 7, which shows the day-5 forecast temperature error averaged over the entire tropical belt (20°N to 20°S) as a function of pressure for both the old and the new model. An average over 6 forecasts from 14 to 19 June 1999 is shown. It is evident that the warm model bias between 300 and 500 hPa has been largely alleviated and the cold bias above the tropical tropopause has been almost halved.

Precipitation

Figure 8 shows a comparison of the frequency bias of 48-hour forecasts of precipitation over Europe for the entire month of August 1999 with the then operational system (blue) and the new CY21r4 system (red). The frequency bias indicates whether the frequency of a precipitation event larger than a given threshold is well simulated by the model. A value of 1 indicates the correct prediction of the frequency of the event, larger values indicate an overestimation and smaller numbers indicate an underestimation. The precipitation thresholds chosen are 0.1, 1, 2, 4, 8 and 16 mm/day, respectively. The new forecasting system improves the forecasts for all threshold classes, except for those larger than 8 mm/day. Both the overestimation in all classes up to 2 mm/day and the underestimation in the class above 16 mm/day are reduced with the new system.

Summary of improvements to the forecasting system

The entire set of changes described in the above sections has been combined to form the new cycle CY21r4 of the IFS, and extensive testing of its forecast performance has been carried out. Figure 9 shows the 500 hPa geopotential height anomaly correlation scores averaged over 131 forecasts from 6 May to 26 September 1999 for Europe, the Northern Hemisphere and the Southern Hemisphere. The scores indicate a substantial improvement of the forecasting system especially over Europe and the Northern Hemisphere. This result is particularly encouraging, since the forecast performance of the previous operational system over that period included spells of rather poor forecast skill in May and August. Other benefits of the new forecasting system can be summarised as:

- ◆ better use of observations in the data assimilation system through improved background error statistics and satellite bias correction schemes;
- ◆ better analysis of low-level winds through the use of SSM/I wind-speed retrievals;
- ◆ more realistic orography fields;
- ◆ improved forecasts of low-level cloudiness, low-level relative humidity, precipitation and 10 m wind speed;
- ◆ improved vertical temperature structure in the tropics.




Figure 7: Vertical profiles of the day-5 temperature errors for forecasts from 14 to 19 June averaged of the tropical belt (20°N to 20°S) for the previous operational system (black line-OPS) and the CY21r4 system (red line-New).




Figure 8: Frequency bias for 48-hour forecasts of precipitation categories larger than 0.1, 1, 2, 4, 8 and 16 mm/day averaged over Europe for the previous operational system (blue line) and for the CY21r4 system (red line).




Figure 9: Anomaly correlation scores of the 500 hPa geopotential height for 131 forecasts from May to September 1999 carried out with the previous operational system (blue line) and the CY21r4 system (red line). The areas shown are Europe (top), Northern Hemisphere (middle) and Southern Hemisphere (bottom).

References

- Buizza, R., Miller, M. and Palmer, T.N.** (1999) Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **125**, 2887-2908.
- Buizza, R. and Palmer, T.N.** (1999) Ensemble data assimilation. Proceedings of the AMS 17th conference on weather analysis and forecasting, 13-17 September 1999, Denver, 231-234.
- Jakob, C. and Klein, S.A.** (1999) The role of vertically varying cloud fraction in the parametrization of microphysical processes in the ECMWF model. *Q. J. R. Meteorol. Soc.*, **125**, 941-965.
- Jakob, C. and Klein, S.A.** (1999) A parametrization of the effects of cloud and precipitation overlap for use in general circulation models. ECMWF Tech. Memo. 289.
- Teixeira, T.** (1999) The impact of increased boundary layer vertical resolution on the ECMWF forecast system. ECMWF Tech. Memo. 268.
- Untch, A., Simmons, A. and colleagues** (1998) Increased stratospheric resolution in the ECMWF forecasting system. ECMWF Newsletter 82, Winter 98/99, 2-8.

C. Jakob, E. Andersson, A. Beljaars, R. Buizza, M. Fisher, E. Gérard, A. Ghelli, P. Janssen, G. Kelly, A.P. McNally, M. Miller, A. Simmons, J. Teixeira and P. Viterbo

Verifying precipitation forecasts using upscaled observations

Verification of precipitation forecasts against SYNOP observations might seem a pretty straightforward procedure, but providing an interpretation of the results is a very difficult task. The model predicts precipitation fluxes over areas of the order of 60 x 60 km², while SYNOP stations report values from rain gauges that typically represent areas of a few tens of cm². It is not simple to bridge this gap, as it cannot be claimed that the spectrum of precipitation field drops to zero for scales below 60 km.

A well-posed verification problem is formulated when several local observations of the precipitation flux inside each model grid box are considered. These observations can be obtained from a high-resolution observing network, or from calibrated radar data. A simple hypothesis is to consider each of these observations equally likely to

contribute to the precipitation flux in a grid box, as there is usually no information on how large the area that each of them represent is. The 'upscaling' as defined in this study consists of averaging all the observations in a model grid box. The super-observations resulting from this procedure are expected to be representative of the grid box.

In this paper the two approaches (using local and upscaled observations) are compared using both error maps and classical threat scores. Local observations are restricted to those received in real time at ECMWF (SYNOP data available on the GTS) while upscaled observations come from a special dataset obtained from Météo-France.

Météo-France has made available climatological precipitation data for 1997, following a request from ECMWF. The