TECHNICAL MEMORANDUM

204

Recent studies of semi-Lagrangian advection at ECMWF

M. Hortal

Research Department

September 1994

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

ABSTRACT

The excess eddy kinetic energy present in the first high resolution semi-Lagrangian version of the ECMWF forecast model has been found to be due to overshooting in the cubic interpolation used for the semi-Lagrangian advection. A quasi-monotone interpolation scheme has been incorporated to avoid this problem.

Low level noise seen in almost all conditions over sea near steep orography has been found to be related not to the size of the orographic ripples in physical space but to the shape of the short-wave tail of the spectrum of the orography. A quick fix for the problem will be to apply a high order diffusion operator to the orography when it is computed from the original dataset.

A "linear" Gaussian grid is presented which does not eliminate the aliasing in the quadratic terms of the equations. An Eulerian treatment of advection with this grid leads to computational instability after a while but with the semi-Lagrangian advection scheme no negative impact is found. This grid reduces drastically the overshooting and undershooting of the fields that are spectrally fitted, therefore supplying the physical parametrizations a set of more realistic (almost free from Gibb's ripples) fields.

1. INTRODUCTION

The semi-Lagrangian approach to integrating NWP models has had considerable development during the past few years. At ECMWF a semi-Lagrangian advection scheme was incorporated into the operational global forecast model in September 1991 following the approach of *Ritchie* (1987). At the same time there was a resolution increase in both the horizontal (where the spectral resolution was increased from T106 to T213) and the vertical (where the number of levels was increased from 19 to 31).

Although the detail obtained from this higher resolution forecast was much better than that obtained from the previous model (mainly at ranges shorter than 5 days), the users noticed a marked increase in the inconsistency between consecutive forecasts verifying at the same time.

Much work was put into understanding this characteristic of the new model and as a result, the semi-Lagrangian treatment was changed in August 1992. The new scheme is non-interpolating in the vertical (*Ritchie et al*, 1994). It gives a better eddy kinetic energy evolution throughout the forecast and a smaller inconsistency between forecasts, although the reason for this improvement was not quite clear (at the same time as this change in the operational model [cycle 43] some other changes were introduced in the physical parametrizations of clouds, radiation, sea-ice temperature and convection which also helped in improving the quality of the forecasts. The changes are documented in the Research Department memorandum R60.6.1/193/NA).

The present paper presents further research into the problem and suggests possible cures.

Another problem found in the high resolution model is noise in the vertical velocity in the low levels of the model in the vicinity of even quite small orographic features. It is present in almost all forecasts, independently of the synoptic situation, and is clearly seen over the sea near Taiwan or in a circular pattern around Hawaii. This noise is not directly related to the semi-Lagrangian advection, but it becomes very marked due to the long time-step used with the semi-Lagrangian scheme. An attempt was made to eliminate this noise by increasing the horizontal diffusion in the model, but the magnitude of the noise in fact increased with this change. The noise was also present when the Eulerian scheme was used with a large time-step stabilized by extra diffusion at upper levels (C Temperton private communication). Although this noise is largest at T213 resolution, it can also be detected in maps of vertical velocity at T106 resolution. At T213 it sometimes produces unrealistic rain bands near steep orography which can be mis-interpreted by the users of the forecast as being a consequence of real mountain waves.

In section 2 a brief description of the semi-Lagrangian treatment of advection is given, together with the idea behind the non-interpolating in the vertical semi-Lagrangian scheme used operationally at ECMWF.

A comparison between the two schemes which have been used operationally is outlined in section 3. Some one-dimensional tests are presented in section 4 to illustrate in a simple way the performance of the techniques tested during this research. The implementation of these techniques in 2 and 3 dimensions are described in sections 5 and 6. In section 7, some of the results obtained using different forms of interpolation to the departure point in the standard semi-Lagrangian treatment are presented. The results obtained so far with the use of the "linear Gaussian grid" are presented in section 8. Finally a summary of the conclusions and an outline of future work is presented in section 9.

2. A DESCRIPTION OF THE SEMI-LAGRANGIAN SCHEME

The prototype equation for the semi-Lagrangian treatment of advection can be written in the form:

$$\frac{d}{dt}X(r,t) = F(X) \tag{1}$$

The simplest three time-level semi-Lagrangian discretization of this equation is:

$$X(r_G,t+\Delta t) = X(r_G-2v\,\Delta t,t-\Delta t) + F(X(r_G-v\,\Delta t,t))\,2\,\Delta t \tag{2}$$

where r_G stands for the arrival point of a parcel of air at the next time step (one of the points of the Gaussian grid), $r_G-2v\Delta t$ is the departure point of this parcel at the previous time step and $r_G-v\Delta t$ is the mid-point of the trajectory. This discretization will be referred to as the interpolating semi-Lagrangian scheme.

On the other hand, equation (1) can be restated as:

$$\frac{d_H}{dt}X(r,t) + (\dot{\eta}^* + \dot{\eta}')\frac{\partial X}{\partial \eta} = F(X) \tag{3}$$

where d_H stands for the advection along a "horizontal" surface (constant value of the vertical coordinate η) and the vertical advection has been split into a part corresponding to the vertical velocity $\dot{\eta}^*$ the air parcel would have if it had departed from the nearest full vertical level to the real departure point, plus the part corresponding to a residual vertical velocity $\dot{\eta}'(=\dot{\eta}-\dot{\eta}^*)$.

If we treat the first part of the vertical advection in (3) in semi-Lagrangian fashion and include the residual vertical advection into the rhs of the equation, we have the scheme which will be referred to as non-interpolating in the vertical semi-Lagrangian scheme, or SLNI.

For more information about the implementation of the semi-Lagrangian scheme in the ECMWF model, see *Ritchie et al* (1994).

3. COMPARISON BETWEEN THE TWO SEMI-LAGRANGIAN SCHEMES

Most of the experimentation discussed in this paper consisted of series of six runs performed at a resolution of T106/L31 (triangular spectral truncation at wavenumber 106 in the horizontal and 31 levels in the vertical using the hybrid coordinate described in *Simmons & Burridge*, 1981), starting from initial data interpolated from the operational analyses corresponding to the 15th of each second month since April 1993. Each series used a different way of interpolating the fields in the semi-Lagrangian treatment. So far only a relatively small number of runs have been performed at the resolution of T213/L31 with the same 31 levels in the vertical. The results are much clearer in the higher resolution experimentation and therefore we present here mainly the results of the high resolution experiments. For the schemes which showed better potential we made a total of 12 experiments in each series, covering the period from March '93 to February '94. The model is the operational ECMWF model with full physics included.

Fig 1(a) and 1(b) show some mean scores of a series of 6 forecasts run from initial data distributed throughout the year at resolution of T106 and 31 vertical levels, using both the non-interpolating in the vertical semi-Lagrangian (SLNI) scheme and the fully interpolating semi-Lagrangian scheme. Figs 2(a) and 2(b) show the same type of scores for a series of 12 forecasts at resolution of T213. One can see that the interpolating semi-Lagrangian scheme performs worse than the SLNI at least over the Northern Hemisphere and over Europe, particularly in terms of root mean square error, and particularly at the higher horizontal resolutions.

Fig 3 presents the mean eddy kinetic energy of the T213/L31 runs. The larger increase of this quantity during the forecast for the interpolating scheme, mainly in the Southern Hemisphere, is a characteristic of all cases studied so far using either the IFS model (or Integrated Forecasting System which is the present operational forecast model at ECMWF since March 1994) or the old operational model of ECMWF (the semi-Lagrangian formulation in the two models differs only in the accuracy of the treatment of the spherical geometry). This excess of eddy kinetic energy is thought to be related to the excessive inconsistency between successive forecasts noticed by the users of the operational products when the fully interpolating semi-Lagrangian scheme was the operational version.

The plotted eddy kinetic energy is a mean over wave numbers up to 42 and therefore the value of the horizontal diffusion applied to the model has a negligible direct impact on it. The corresponding curve for the analysis includes wavenumbers up to 63 in the standard verification package and therefore is expected to be always above the forecast ones.

One hypothesis to explain the lower increase of eddy kinetic energy in the SLNI version was the occurrence of structures in the vertical containing a large $2\Delta\eta$ noise in the non-interpolating-in-the-vertical scheme but

not in the interpolating scheme. These structures could make the vertical diffusion work much harder, so producing a decrease of the eddy kinetic energy.

To test this hypothesis, a series of experiments was run without vertical diffusion in the free atmosphere, both with the SLNI and with the interpolating scheme. The difference in eddy kinetic energy with the corresponding runs with vertical diffusion was very similar in both schemes, indicating that the effect of the vertical diffusion is not very different for the two semi-Lagrangian versions.

Global diagnostics indicate that the main difference in the contribution to the total kinetic energy between the SLNI runs and the interpolating runs comes from the pressure gradient term. This could indicate that the fully interpolating scheme has a larger gravity wave activity than the SLNI scheme.

In order to try to understand how this higher activity could arise from the interpolating scheme, a series of simple tests were made applying the semi-Lagrangian treatment to the one-dimensional linear advection equation.

4. LINEAR ONE-DIMENSIONAL TESTS

In order to understand the "production" of eddy kinetic energy by the semi-Lagrangian advection process, a unidimensional test was made in which a sharp initial field was advected in a grid-point calculation by a constant wind such that the CFL number was 2.5. In this way, the interpolations to the departure point in the semi-Lagrangian procedure were always made to the centre point between two grid points so that one could assess the effect of the interpolations.

Fig 4(a) shows the initial field while Fig 4(b) shows the result after 1000 time-steps using cubic interpolation to the departure point of each trajectory. The overshooting and undershooting produced by the high order interpolation are clearly visible. In this simple case of the linear advection equation, the overshooting and undershooting are simply advected and possibly amplified in future time steps. In a full PE model, they represent local perturbations, presumably not in quasi-geostrophic equilibrium, which will produce gravity waves and therefore a higher level of eddy kinetic energy in the model.

According to this interpretation, the production of excess kinetic energy could arise, not only from the vertical interpolation but also from the horizontal cubic interpolation at the places where the interpolated fields are sharp. The fact that this is most noticeable in the fully interpolating scheme could be understood if the overshooting produced by the horizontal interpolation were further amplified by the cubic vertical interpolation.

In a second test, the spectral method used in the forecast model of ECMWF was simulated by transforming to Fourier space and back to grid-point space after each advection step. The relationship between the Fourier truncation M and the number of grid-points N is the same as used in the Gaussian grid of the spectral method, namely

$$M = (N-1)/3.$$

This truncation produces spectral ripples in the same manner as spectral ripples are produced in the grid-point fields of a spectral forecast model. As can be seen in Fig 4(c), the combined presence of spectral ripples and cubic interpolation very slightly worsens the overshooting and undershooting in the advected field.

In order to reduce the problem, a quasi-monotone cubic interpolation was used as proposed by *Bermejo and Staniforth* (1992). The result of this procedure is shown in Fig 4(d). The overshooting and undershooting have been substantially reduced. A further improvement can be achieved by increasing the truncation limit in the Fourier space to

$$M = (N-1)/2.$$

This grid does not eliminate the aliasing in quadratic terms and therefore could not be used in connection with an Eulerian treatment of non-linear advection. However, as discussed by *Côté & Staniforth* (1988) it can be used in connection with the semi-Lagrangian advection procedure because the non-linearity of the advection process is not handled directly as the product of two spectrally represented variables but as a linear combination of grid-points which is the interpolation to the departure points of the trajectories. The result of using this grid is to avoid spectral ripples in the transforms to Fourier space and back. If the quasi-monotone cubic interpolation is then used in the semi-Lagrangian procedure, we eliminate completely the overshooting and undershooting, with the result that the amplitude of the advected field is never larger than in the initial field, Fig 4(e). This constitutes the basis for considering the "linear Gaussian grid" for a spectral semi-Lagrangian model as discussed later in section 8 of this paper.

5. QUASI-MONOTONE CUBIC INTERPOLATION IN A THREE-DIMENSIONAL MODEL

Bermejo and Staniforth (1992) describe the Quasi-Monotone Semi-Lagrangian (QMSL) algorithm for bidimensional interpolation as follows:

Let $(U1, U2, U3, U4)$ be the values of the field to be interpolated at the vertices of the grid element in which the departure point of a trajectory is located. Then we set

$$U^+ = \max(U1, U2, U3, U4)$$
$$U^- = \min(U1, U2, U3, U4)$$

The steps of the QMSL algorithm are then:

1) Compute the departure point and identify the element of the grid where such a point is located.

2) Compute the high-order approximation U_H employing any standard interpolation of order higher than 1.

3) Evaluate U^+ and U^-.

4) Set the estimated departure value U of the function

   $U = U^+$ if $U_H > U^+$

   $U = U^-$ if $U_H < U^-$

   $U = U_H$ otherwise.

We have applied this scheme to the normal bi-cubic Lagrange interpolation used in the horizontal in the semi-Lagrangian version of the ECMWF forecast model and found a small decrease in eddy kinetic energy of the model.

A larger benefit was found when the clipping procedure described in point 4) above was applied at every one-dimensional interpolation. The applied procedure is therefore as follows (refer to the diagram below where A-L are grid points; note that in a reduced Gaussian grid the grid points are not necessarily aligned in the latitudinal direction):

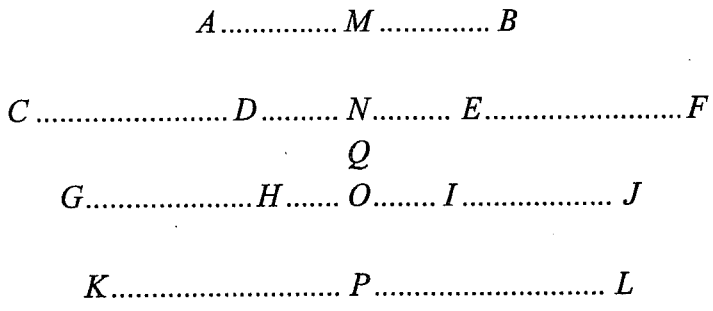

1) Compute the departure point (Q) of each trajectory arriving at the grid points of a certain row of latitude.

2) Do a cubic Lagrange interpolation to the departure longitude at each of the rows north and south of the departure latitude (point N using grid-points C, D, E and F and point O using grid-points G, H, I and J). In both interpolations, if the interpolated value is larger than both the centre values used for the interpolation (points D and E for interpolation to point N and points H and I for point O), reset the interpolated value to the larger of the two of them, with a similar treatment if the interpolated value is smaller than both the adjacent values.

3) Do a linear interpolation to the departure longitude at the two next-nearest neighbouring rows from the departure point (points M and P) (this procedure is adopted for reasons of economy in the ECMWF model for horizontal interpolation).

4) Do cubic Lagrange interpolation in latitude to the departure point Q using the four interpolated values (at M, N, O and P) computed in points 1)-3) and clip the resulting value to stay within the range of the two centre points (N and O) used for the interpolation.

5) For three-dimensional interpolations, do steps 1)-4) for the two nearest levels to the vertical departure position and do bilinear interpolation at the next-nearest levels. Finally do a cubic interpolation in the vertical.

In the vertical interpolation, we could use again the QMSL algorithm but this was found to produce a large positive error in the temperature field in the tropical stratosphere due to the interpolations to points close to the tropopause. Clipping the interpolated values at these positions produces a smoothing of the tropopause (positive temperature errors) which is then advected upwards due to the predominantly ascending vertical velocity in the tropical stratosphere. One can instead blend cubic and linear interpolations in a different way, by using both cubic and linear interpolations in the vertical and taking a linear combination of both as interpolated value. The best blending in terms of mean zonally averaged temperature errors seems to be 1.3 times the cubic interpolated value minus 0.3 times the linearly interpolated value. This in fact produces an interpolation which is sharper than the cubic Lagrange interpolation, giving then a better representation of the tropopause and smaller errors in the stratosphere. Nevertheless, the runs made with this kind of vertical interpolation for the temperature field showed a lower level of skill in terms of objective verification scores.

6. THE "LINEAR GAUSSIAN GRID"

In a further series of experiments, we adopted the grid named "reduced-resolution Gaussian grid" by *Côté and Staniforth* (1988). This grid is defined as the minimum sized grid required to transform from spectral to grid point space and back again without loss of information.

The number of points per row in this grid will be such that the Fast Fourier Transforms are exact, i.e. $(M+1)*2$, where M is the spectral truncation, and the corresponding number of rows in the Gaussian grid will be $(M+1)$. This grid would therefore have $(2/3)^2$ the number of points than the grid usually chosen has, and should represent exactly only the linear terms. Our term for this grid is chosen to avoid confusion with the term "reduced Gaussian grid" used by *Hortal and Simmons* (1991) which entails a reduction in the number of points used at each latitude row as one approaches the pole, in order to keep the resolution in grid-point space more uniform than in the normal Gaussian grid. The grid actually used in the experiments reported in this paper is a "reduced linear grid" in the sense that it uses $(M+1)$ rows with $(M+1)*2$ points

at the rows nearest to the equator and then reduces this number as one approaches the poles, in the same way as described in *Hortal and Simmons* (1991).

With triangular truncation, there still a larger number of degrees of freedom for real-valued fields in grid-point space than in spectral space. Thus the transformation to spectral space and back to grid-point space is not exact even for linear terms. Even so, the problems with spectral ripples in grid-point space, if not completely eliminated, should be reduced by the use of this "linear grid". This point is illustrated in Fig 5, where we show the spectrally fitted orography near the Andes using the same 160-latitudes Gaussian grid but using (Fig 5a) a triangular truncation of T106 in spectral space (the usual resolution for this Gaussian grid) or (Fig 5b) a triangular truncation of T159 (for which the "linear Gaussian grid" coincides with the usual Gaussian grid for T106). The negative values (dashed lines in the figures) caused by the spectral fitting have been much reduced. This can be seen more clearly in Fig 5(c) where we show a cross-section of the surface geopotential along the horizontal line drawn in the upper left hand corner of panels (a) and (b). Here the full line corresponds to the orography fitted at spectral resolution T106 while the dotted line corresponds to the orography fitted at T159; also shown is the cross-section (dash line) corresponding to the "diffused" orography proposed in section 8. The same reduction of negative values can be seen in the orography over Europe, shown in panels (d) and (e) of the same figure.

7. EXPERIMENTS USING THE STANDARD GAUSSIAN GRID

A series of 6 experiments were run at resolution T213 using the standard SLNI, from initial data at 12:00 UT the 15th of each second month from April 93 to February 94. This series is used as control for comparison with the other schemes described in this paper.

A second series used the fully interpolating scheme, with cubic Lagrange interpolation for the terms evaluated at the departure point of each trajectory, and linear interpolation for the terms evaluated at the centre of the trajectories. These are the same interpolations as those used in the non-interpolating scheme (SLNI) except that in the SLNI scheme, the interpolations to the departure point of the trajectories are always bi-dimensional.

In the third series of experiments, using also the fully interpolating scheme, the horizontal interpolations were performed by means of the quasi-monotone scheme described previously, this series will be called SLIQ.

In all three series, the right hand sides of the thermodynamic and the continuity equations were evaluated at the centre of the trajectory while the right hand sides of the momentum equations were evaluated as the mean of the values interpolated at the origin of the trajectory and at the end of the trajectory. This

configuration has been found to give better scores and smaller level of noise than any other tested previously.

The fourth series also used quasi-monotone cubic interpolations in the horizontal and standard cubic interpolation in the vertical for the terms corresponding to time $t-\Delta t$. The right hand sides of all the equations (the terms corresponding to time t) were evaluated as an average along the trajectory but the interpolation of these terms at the origin point made with trilinear interpolation instead of the standard procedure; the latter consists of adding these terms those corresponding to time t-Δt and performing cubic interpolation for the sum. This series will be called HQM3.

Fig 6(a) shows some mean scores for the four series for geopotential, while Fig 6(b) shows scores for vector wind and temperature.

In the Northern Hemisphere and in Europe, where the non-interpolating scheme performed better than the interpolating scheme, the usage of quasi-monotone interpolation in the horizontal brings the scores to the same level as those for the non-interpolating scheme. The averaging of the right hand side of the equations along the trajectory with linear interpolation at the departure point for these terms improves further the scores at the later stages of the forecast. In the Southern Hemisphere, where the interpolating scheme performed better than the non-interpolating, the use of quasi-monotone interpolation also improves the scores.

One has to be careful in interpreting some of the scores, particularly the root mean square error at low levels because the vertical profile of the fields near the surface is rather different in the two vertical treatments and therefore the post-processing, which uses extrapolation downward from the second level above the surface, favours the operational scheme (the SLNI scheme) over the experimental ones (interpolating in the vertical).

Fig 7 (a-d) shows the mean zonal error of the temperature field averaged over the range of the forecasts (10 days). The most striking characteristic of this error is the large positive error in the stratosphere in the SLNI runs. As mentioned earlier, this error is increased if we use quasi-monotone interpolation of the temperature in the vertically interpolating semi-Lagrangian scheme. The error may therefore be due to a smoothing of the tropopause in the vertical advection computation. The error is present to a much smaller extent in all the interpolating tests, independent of the kind of interpolation used in the horizontal. Panels (e-h) of Fig 7 show the zonal mean error of the zonal wind component. The SLNI scheme shows the largest error in the strength of the polar jet in the Northern Hemisphere and the HQM3 scheme using quasi-monotone interpolation in the horizontal and averaging of the RHS of the equations along the trajectory shows the smallest overall error in this field, both in the northern and in the Southern Hemispheres.

Fig 8 shows the zonal mean error of four series of experiments performed at the lower resolution of T106 (the scores in this case are virtually identical for the four series and will not be presented here). The positive error in the temperature in the tropical stratosphere is of the same order of magnitude as in the T213 experiments, but the reduction achieved by the use of the quasi-monotone interpolation in the horizontal is smaller. This error can be further reduced if we use also quasi-monotone interpolation of the wind fields in the vertical (series QM2) and still further if the vertical interpolation of the temperature field is "sharpening" as described at the end of section 5 (series QM4). The influence of the different interpolations on the error of the zonal wind is quite small. No definite conclusions could be drawn from this series of experiments until the corresponding T213 experiments can be run.

In order to check the assumption that the horizontal cubic interpolation produces an excess eddy kinetic energy which affects negatively the scores, a series of experiments was run using the SLNI scheme (which uses only horizontal interpolation at the departure points) but with quasi-monotone cubic interpolations (series NIQM). Fig 9(a-b) shows some of the scores from this series, compared with the standard SLNI scheme. All the scores were improved by the quasi-monotone interpolation. The zonal mean temperature error was not changed appreciably and the zonal mean zonal wind error changed only slightly. Also shown in the figure are the scores of a series (QMUV) using the fully interpolating semi-Lagrangian scheme and quasi-monotone interpolation only for the wind fields, in the three dimensions. The corresponding eddy kinetic energy curves are shown if Fig 10. The reduction due to the quasi-monotone interpolation is, as expected, larger in the case of the fully interpolating scheme than in the case of the SLNI scheme, as the interpolations are three-dimensional rather than bi-dimensional.

The most recent series of experiments at T213/L31 resolution uses the quasi-monotone interpolation in the three dimensions for the wind field (and humidity field) and only in the horizontal for the temperature (series QMT3). This will allow an evaluation of the relative importance on the eddy kinetic energy of using the quasi-monotone interpolation for the temperature compared with the wind and in the horizontal compared with the vertical. As can be seen in panels (d-f) of Fig 10, (which corresponds to T106 series), this is the configuration which gives the smallest excess eddy kinetic energy.

The result on the eddy kinetic energy for this series at T213 is shown in Fig 11, together with three other series of T213 experiments. Series QMUV uses quasi-monotone interpolation in the three dimensions for the wind fields only. The reduction in eddy kinetic energy as compared with the SLI scheme is significant but the use of the quasi-monotone scheme in the vertical has a rather small influence as can be seen comparing series HQM3 and QMT3. In both these last two series, quasi-monotone interpolation is used for all the fields in the horizontal, the difference being the use of quasi-monotone interpolation in the vertical for the wind in the series QMT3 and not in HQM3. From this comparison it is clear that, in order to reduce

the excess eddy kinetic energy in a semi-Lagrangian model it is more important to apply the quasi-monotone interpolation in the horizontal rather than in the vertical. The fact that the non-interpolating in the vertical scheme has a low excess eddy kinetic energy should be attributed to a compensation of the increase due to horizontal interpolation with the decrease due to the Eulerian (first order accurate) treatment of the vertical advection which smooths the fields.

Is summary, one can say that the use of quasi-monotone interpolations in the horizontal part of the semi-Lagrangian scheme reduces the excess eddy kinetic energy of the model (particularly at high resolution) and improves the scores of the forecasts. The use of quasi-monotone interpolation for the wind fields in the vertical has also a (small) positive effect but its use in the temperature field produces a large positive error in the tropical stratosphere.

A smoothing of the RHS of the equations by averaging its value along the semi-Lagrangian trajectory and using trilinear interpolation for the part corresponding to the departure point has also a beneficial effect on the scores.

8. EXPERIMENTS USING THE "LINEAR GAUSSIAN GRID"

One of the characteristics of the high resolution model used in operations at ECMWF is the low-level noise which can be seen in almost every forecast near steep orography and which can produce unrealistic rain bands in certain circumstances. This noise is also seen in the Eulerian version of T213/L31 if the time step is increased sufficiently and can be illustrated using a map of vertical velocities on the lowest model level near Taiwan [Fig 12 (a); the arrows represent the horizontal wind at the same level].

A possible explanation of this noise could be aliasing in the higher order terms of the equations. As the "linear grid" has different aliasing properties than the standard Gaussian grid, this noise was examined in a set of T159 experiments run with the same grid as the standard T106. In the first of them, the orography was recomputed, doing the spectral fit with T159 resolution in spectral space. The map of vertical velocity at the lowest model level is shown in Fig 12(c). It is clear that the noise has increased substantially in spite of the fact that the amplitude of the orographic ripples is smaller in this fitting of the orography than in the standard T106 fit. In order to check whether this is due to a different representation of the orography or to different treatment of the aliasing, another experiment was run at T159 resolution, using the orography spectrally fitted at T106 resolution. The corresponding map of vertical velocity is shown in Fig 12(b). This map is virtually identical to that of the T106 experiment, indicating that it is the representation of the orography and not the aliasing in the model that is the cause of the noise.

A third experiment was then run, recomputing the orography with T159 spectral fit but applying a reduction in the short-wave tail of its spectrum, by a factor of 10, by means of an 8th order diffusion operator. This

reduced the amplitude of wavenumber 106 by a factor of about 1.4. The resulting vertical velocity is shown in Fig 12(d). It is clear that most of the noise has disappeared and the vertical velocity field looks now much more realistic. This result and the fact that the noise is increased if we increase the horizontal diffusion in the upper air fields leads us to the conclusion that the noise is due to the presence of an orographic spectrum whose short-wave tail is too large compared with the shape of the spectrum of the upper air fields. The horizontal diffusion coefficient applied to the dynamical fields during the forecast for the T159 experiments was chosen such that the damping of wave-number 106 is about half of what it is in the T106 runs. All the three experiments were run with the standard non-interpolating-in-the-vertical semi-Lagrangian scheme but the noise appears not to be affected by the kind of semi-Lagrangian scheme used.

A series of 6 cases has been run at T319/L31 using the SLNI scheme and the "linear reduced" Gaussian grid which coincides with the standard reduced Gaussian grid for resolution T213. The mean scores are shown in Fig 13. There is a need to fine tune the horizontal diffusion coefficients for this resolution, but it is clear that the linear grid does not produce any significant negative effects on the forecast. The cost in terms of CPU time of running at this higher spectral resolution is about 8% higher than running at T213. The horizontal diffusion used was such that the damping of wavenumber 213 is about half of what is in the T213 runs and the size of the time step was the same for both spectral resolutions. Fig 14 shows maps of precipitation to the west of central America between 24 and 48 hours into the forecast, for one of the cases of the series. Panel (a) is the forecast using the linear grid i.e. using a spectral resolution of T319 while panel (b) is the forecast made at spectral resolution T213 and therefore the Gaussian grid corresponds to the standard choice. Panel (c) is the difference (a)-(b). Over sea, the maximum negative differences corresponds to the maxima in precipitation. This means that the precipitation field is smoother at T319 resolution because there are less overshooting of the humidity maxima and undershooting of the humidity minima than in T213 (the humidity field is therefore more realistic). The maximum furthest to the east is produced by the orography. As the orography in the T319 representation is sharper than at T213, this maximum is correspondingly amplified in the forecast at T319.

The study of a linear grid in a semi-Lagrangian spectral model is still in a very early stage. Much more work is needed to compare this grid with the standard Gaussian grid before a definitive statement can be made about the advantages of one over the other.

9. CONCLUSIONS

Some of the problems found since the implementation at ECMWF of a high resolution (T213/L31) semi-Lagrangian model in September '91 have been investigated.

The most important of these problems has been the excess eddy kinetic energy seen in the forecasts which used the fully interpolating semi-Lagrangian scheme. This is thought to be related with increased inconsistency between successive forecasts, which led ECMWF to change the advection scheme in August '92 to a non-interpolating in the vertical semi-Lagrangian scheme.

It has been found that this excess eddy kinetic energy is due mainly to the over/undershooting produced by the cubic interpolation (mainly the horizontal part) at the departure point of the semi-Lagrangian trajectories. The excess eddy kinetic energy is also present in the non-interpolating-in-the-vertical semi-Lagrangian scheme, but to a smaller extent due to the lower dimensionality of the interpolations in this case and to a smoothing of the fields produced by the first order accurate vertical advection.

In order to reduce this spurious production of eddy kinetic energy, we have tested a modified form of the quasi-monotone interpolation proposed by *Bermejo and Staniforth* (1992). The use of this interpolation in the non-interpolating-in-the-vertical scheme reduces somewhat the eddy kinetic energy of the forecasts and improves the objective scores. Using this interpolation in the fully interpolating semi-Lagrangian scheme reduces the eddy kinetic energy by a larger amount, to a level similar to that of the quasi-monotone non-interpolating-in-the-vertical scheme. The objective scores are improved to the level of the operational scheme (non-interpolating in the vertical using cubic interpolation in the horizontal).

A further improvement in the scores at the later stages of the forecast is achieved by averaging the RHS of the equations along the semi-Lagrangian trajectory and using trilinear interpolation for the part of this RHS corresponding to the departure points of the semi-Lagrangian trajectories.

Another problem frequently seen in the operational forecasts has been noise in the vertical velocity at low levels of the model near narrow orographic features. This sometimes produces unrealistic rain bands. An increase of the horizontal diffusion worsens the problem rather than improving it. The low-level noise near orography is due to a mismatch between the shape of the orography spectrum and the spectra of the dynamical fields. The problem can be alleviated by reducing the amplitude of the sort wave tail of the orography spectrum by means of a high order diffusion operator applied during the computation of the orography.

A general problem in spectral models is the presence of Gibbs phenomena due to the spectral fitting with the Gaussian grid normally used to avoid aliasing in the quadratic terms of the equations. To reduce this problem, a different Gaussian grid has been used which does not eliminate aliasing in the quadratic terms. This grid is not suitable for an Eulerian treatment of the advection but does not produce appreciable negative effects when used in conjunction with a semi-Lagrangian scheme. The fields handled by the physical

parametrizations when this grid is used are free to a large extent of the overshooting or undershooting produced by the spectral fitting.

Acknowledgements. My thanks to Adrian Simmons and Tony Hollingsworth for a careful review of the manuscript. To Michael Tiedtke for advocating the use of the "linear grid" in the first place, at least for the treatment of the physical parametrizations, and to Adrian Simmons and Martin Miller for many interesting discussions.

REFERENCES

Bermejo, R and A Staniforth, 1992: The conversion of semi-Lagrangian advection schemes to quasi-monotone schemes. Mon Wea Rev, **120,** 2622-2632.

Côté, J and A Staniforth, 1988: A Two-time-level semi-Lagrangian semi-implicit scheme for spectral models. Mon Wea Rev, **116,** 2003-2012.

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

Ritchie, H, 1987: Mon Wea Rev **115,** 608-619.

Ritchie H, C Temperton, A Simmons, M Hortal, T Davies, D Dent and M Hamrud, 1994: Implementation of the semi-Lagrangian method in a high resolution version of the ECMWF forecast model. Accepted for publication at Mon Wea Rev.

Simmons, A J and D M Burridge, 1981: An energy and angular-momentum conserving vertical finite-difference scheme and hybrid vertical coordinates. Mon Wea Rev, **109,** 758-766.

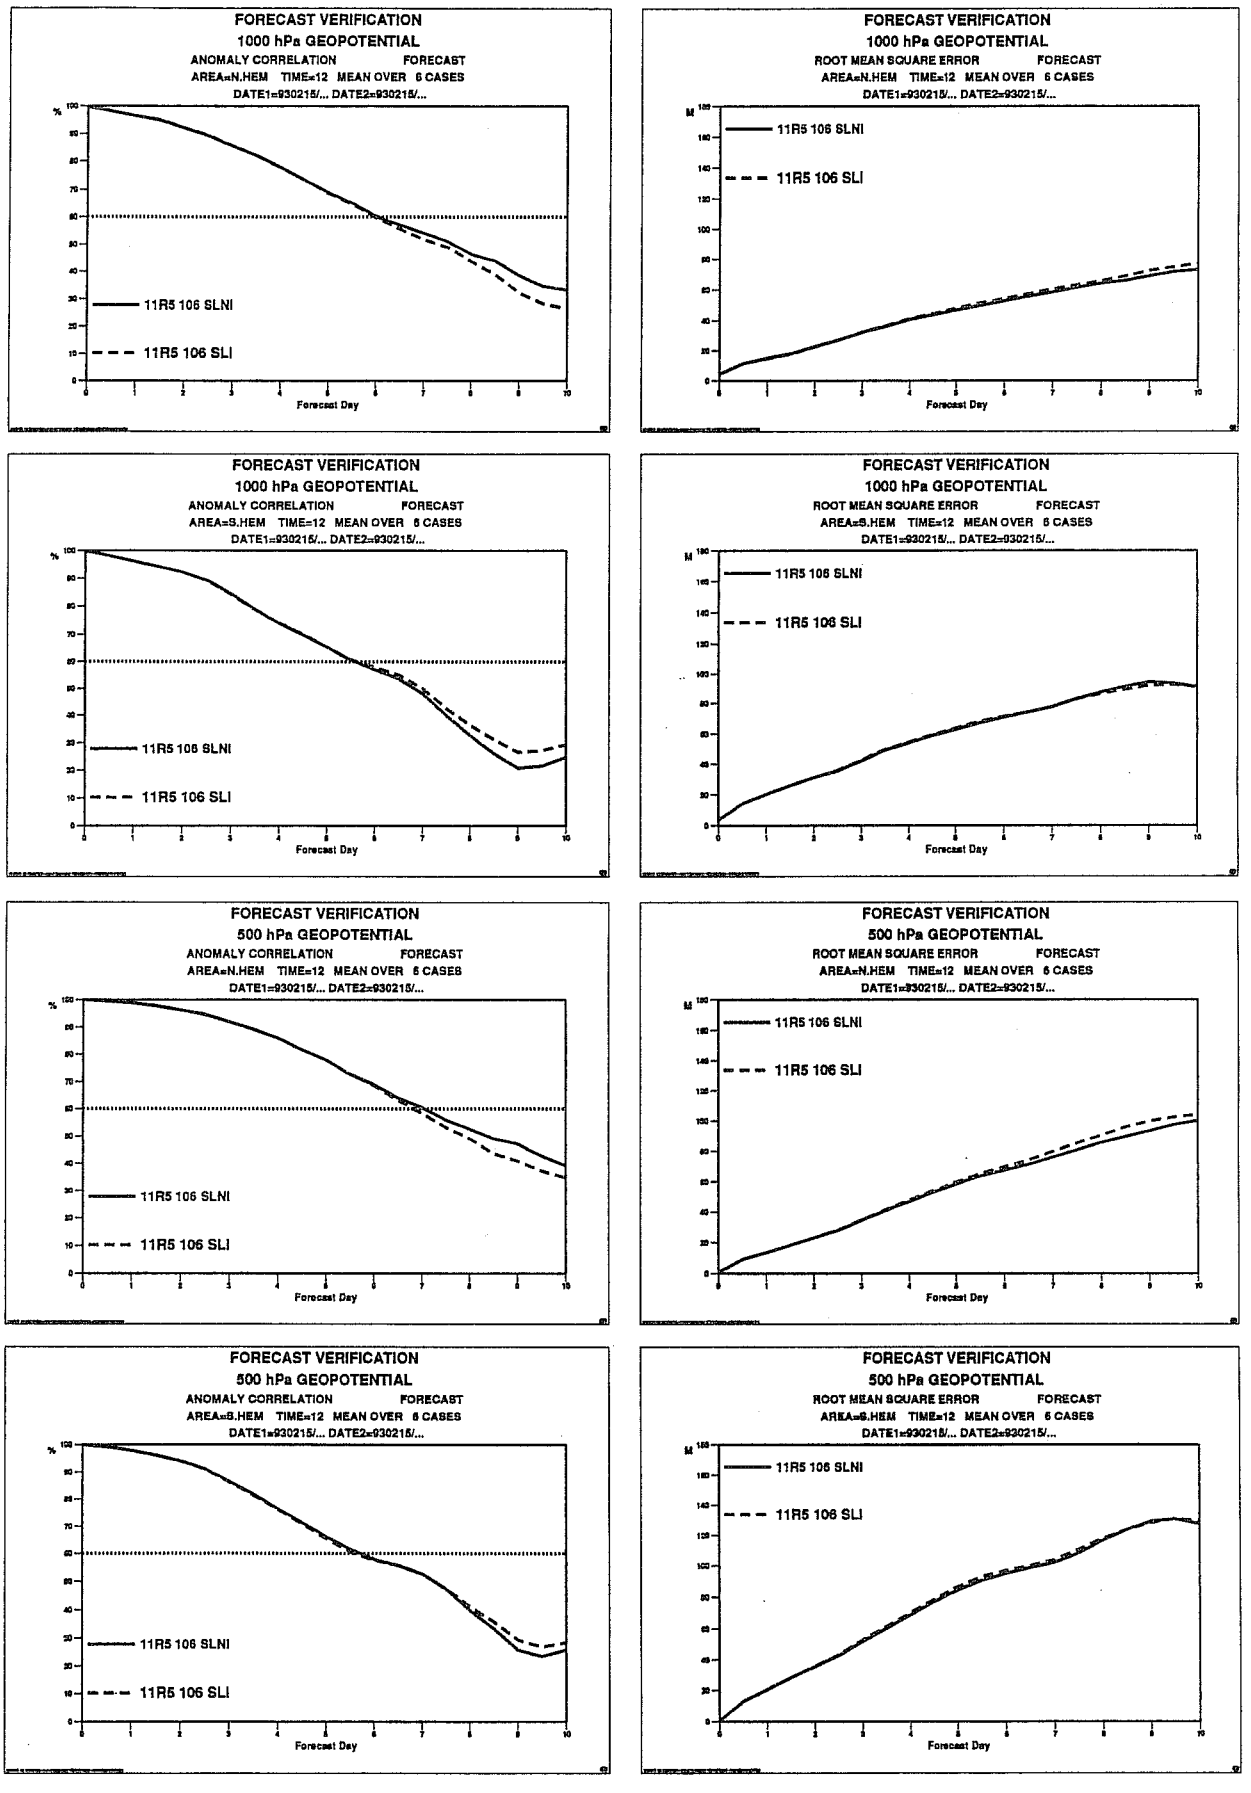

Fig. 1(a): Mean scores of a series of 6 forecasts run from initial data distributed throughout the year at resolution of T106 and 31 vertical levels, using both the non-interpolating in the vertical semi-Lagrangian (SLNI) scheme and the fully interpolating semi-Lagrangian scheme (SLI).

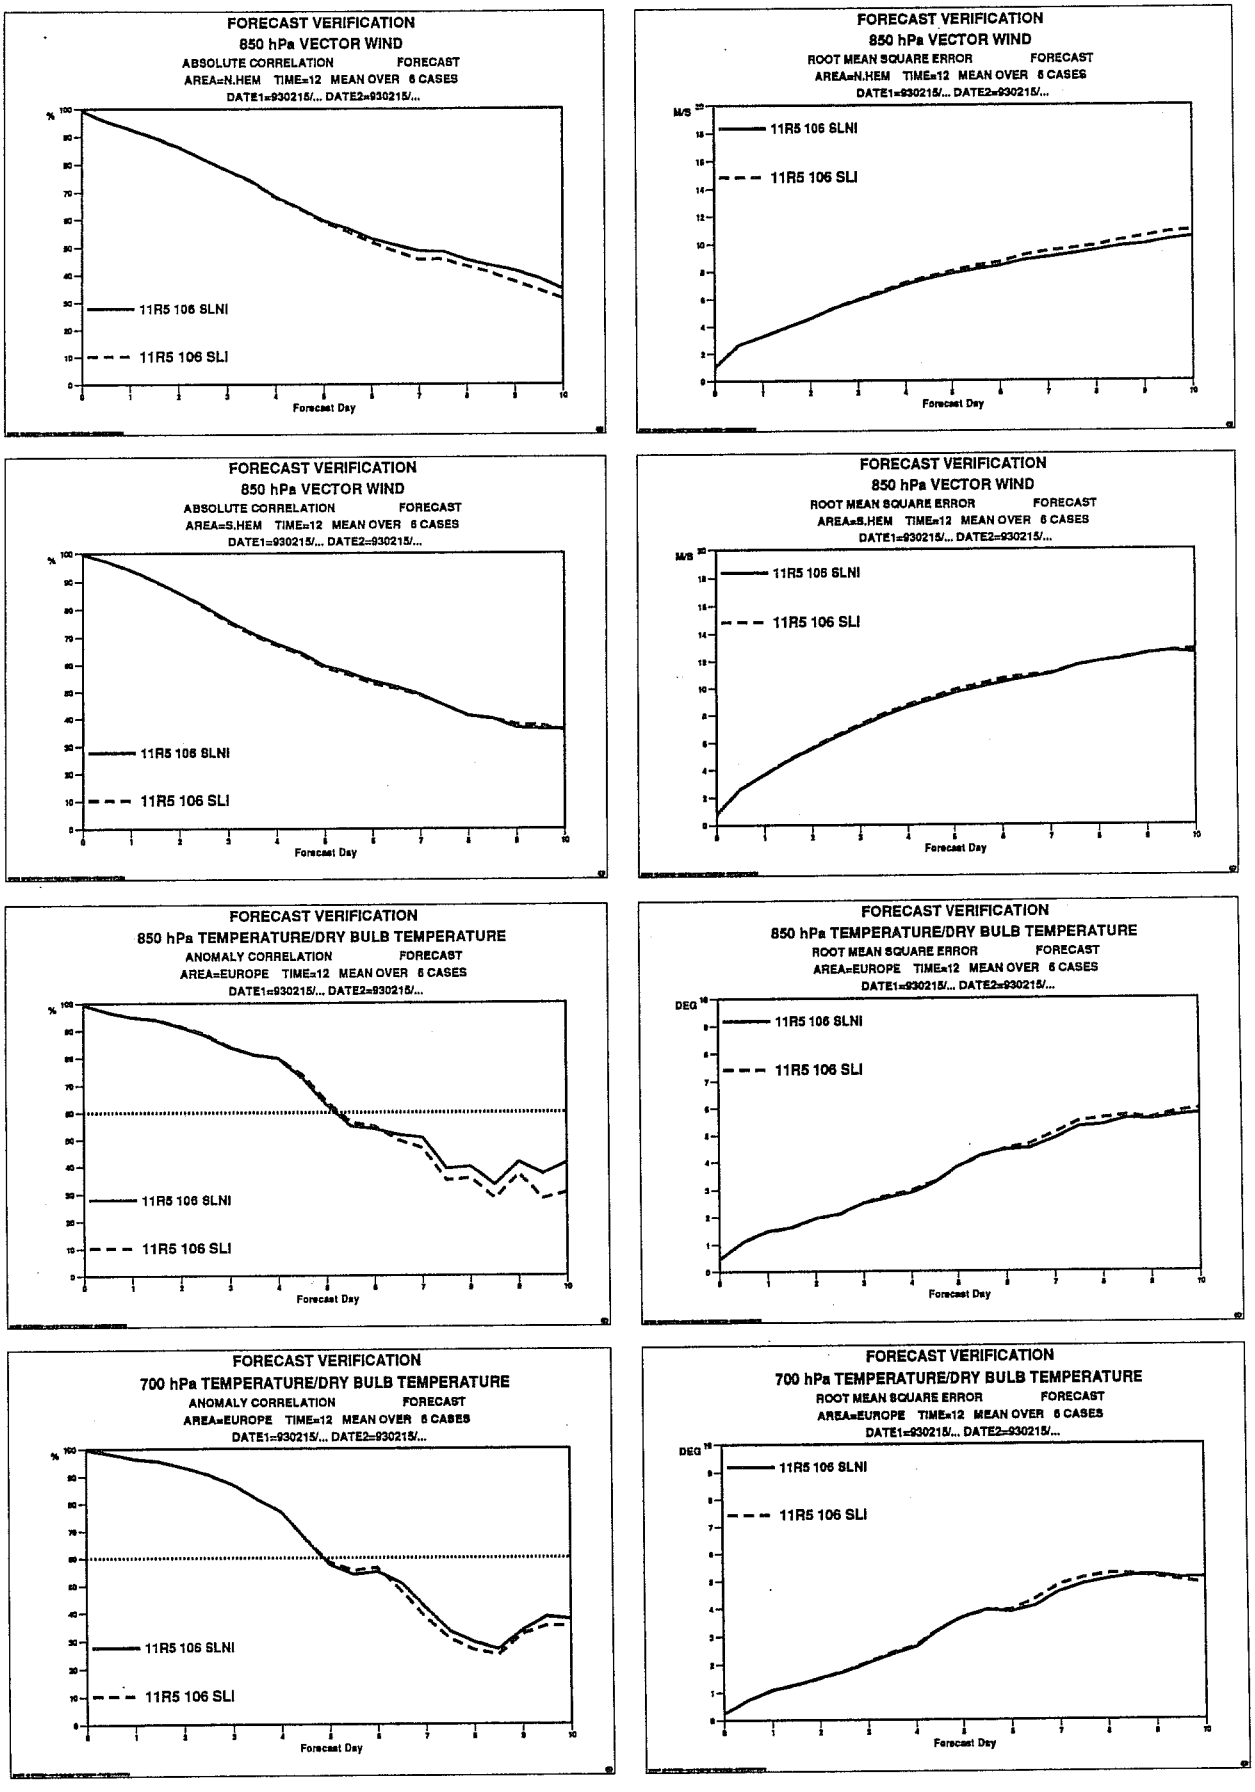

Fig. 1(b): Mean scores of a series of 6 forecasts run from initial data distributed throughout the year at resolution of T106 and 31 vertical levels, using both the non-interpolating in the vertical semi-Lagrangian (SLNI) scheme and the fully interpolating semi-Lagrangian scheme (SLI).

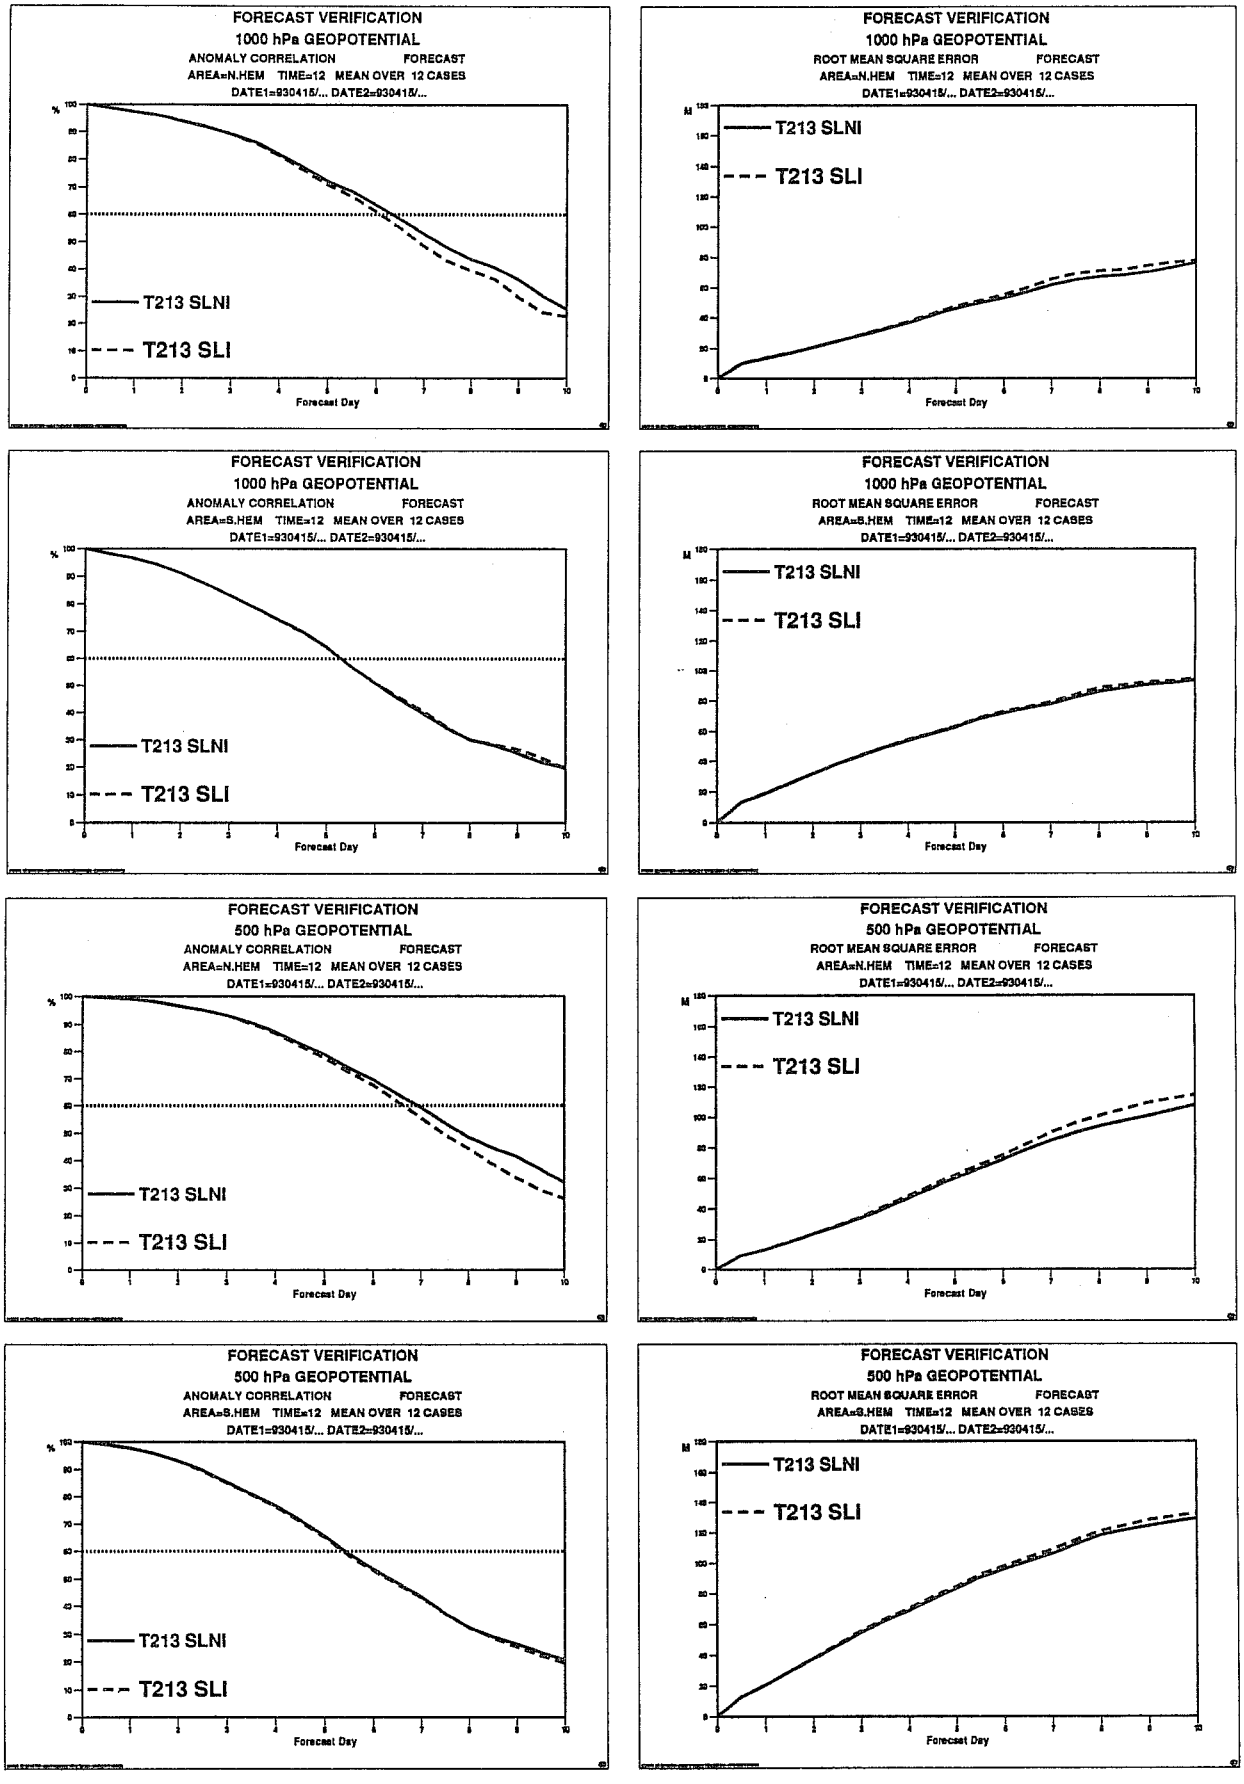

Fig. 2(a): As Fig. 1 but for a series of 12 cases using T213 horizontal resolution.

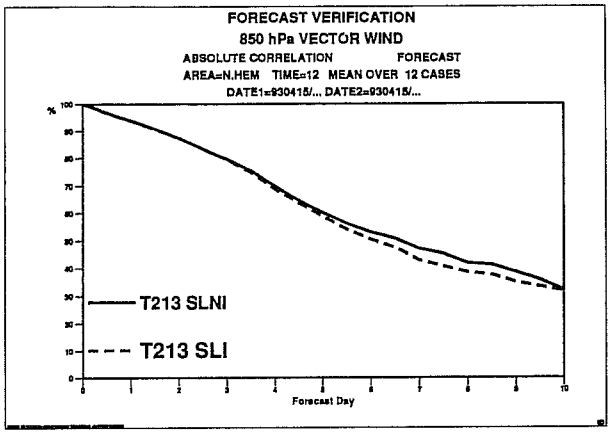

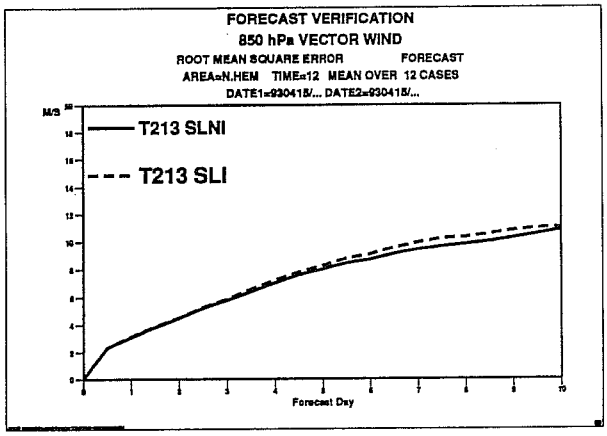

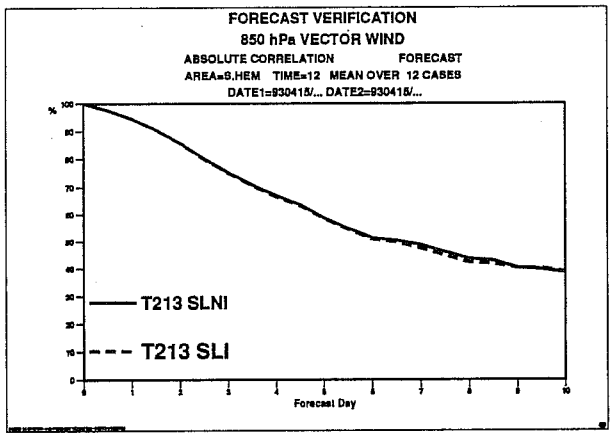

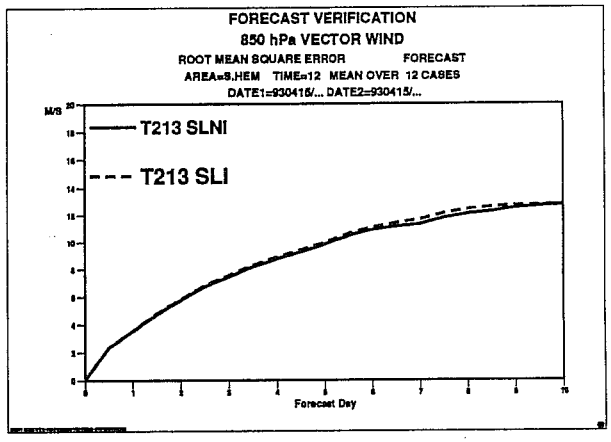

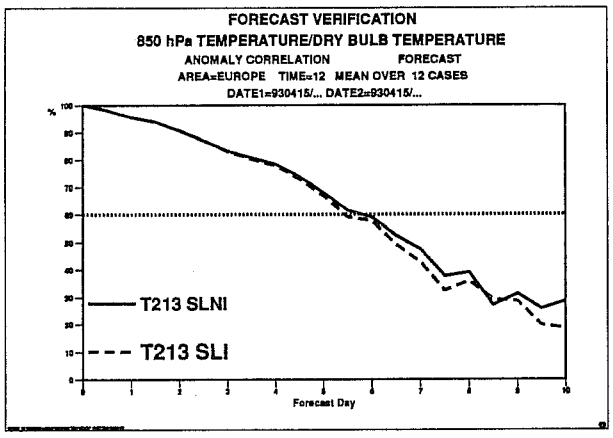

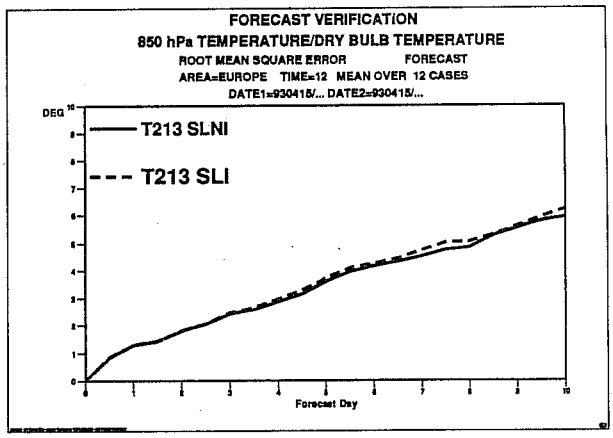

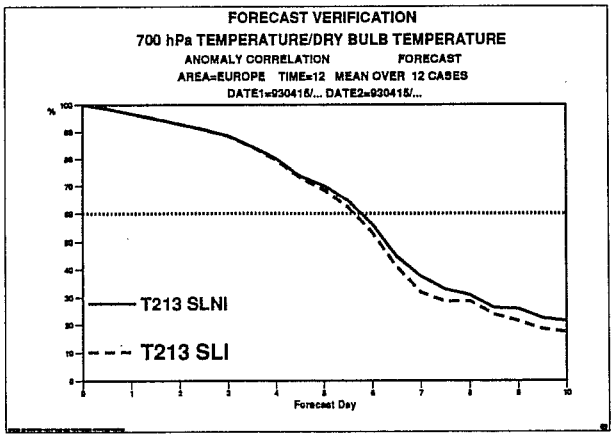

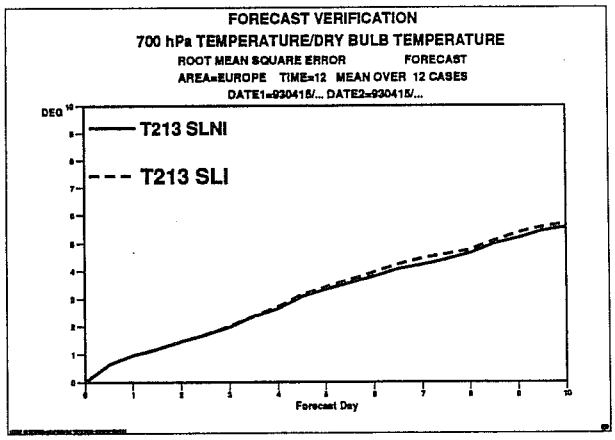

Fig. 2(b): As Fig. 1 but for a series of 12 cases using T213 horizontal resolution.

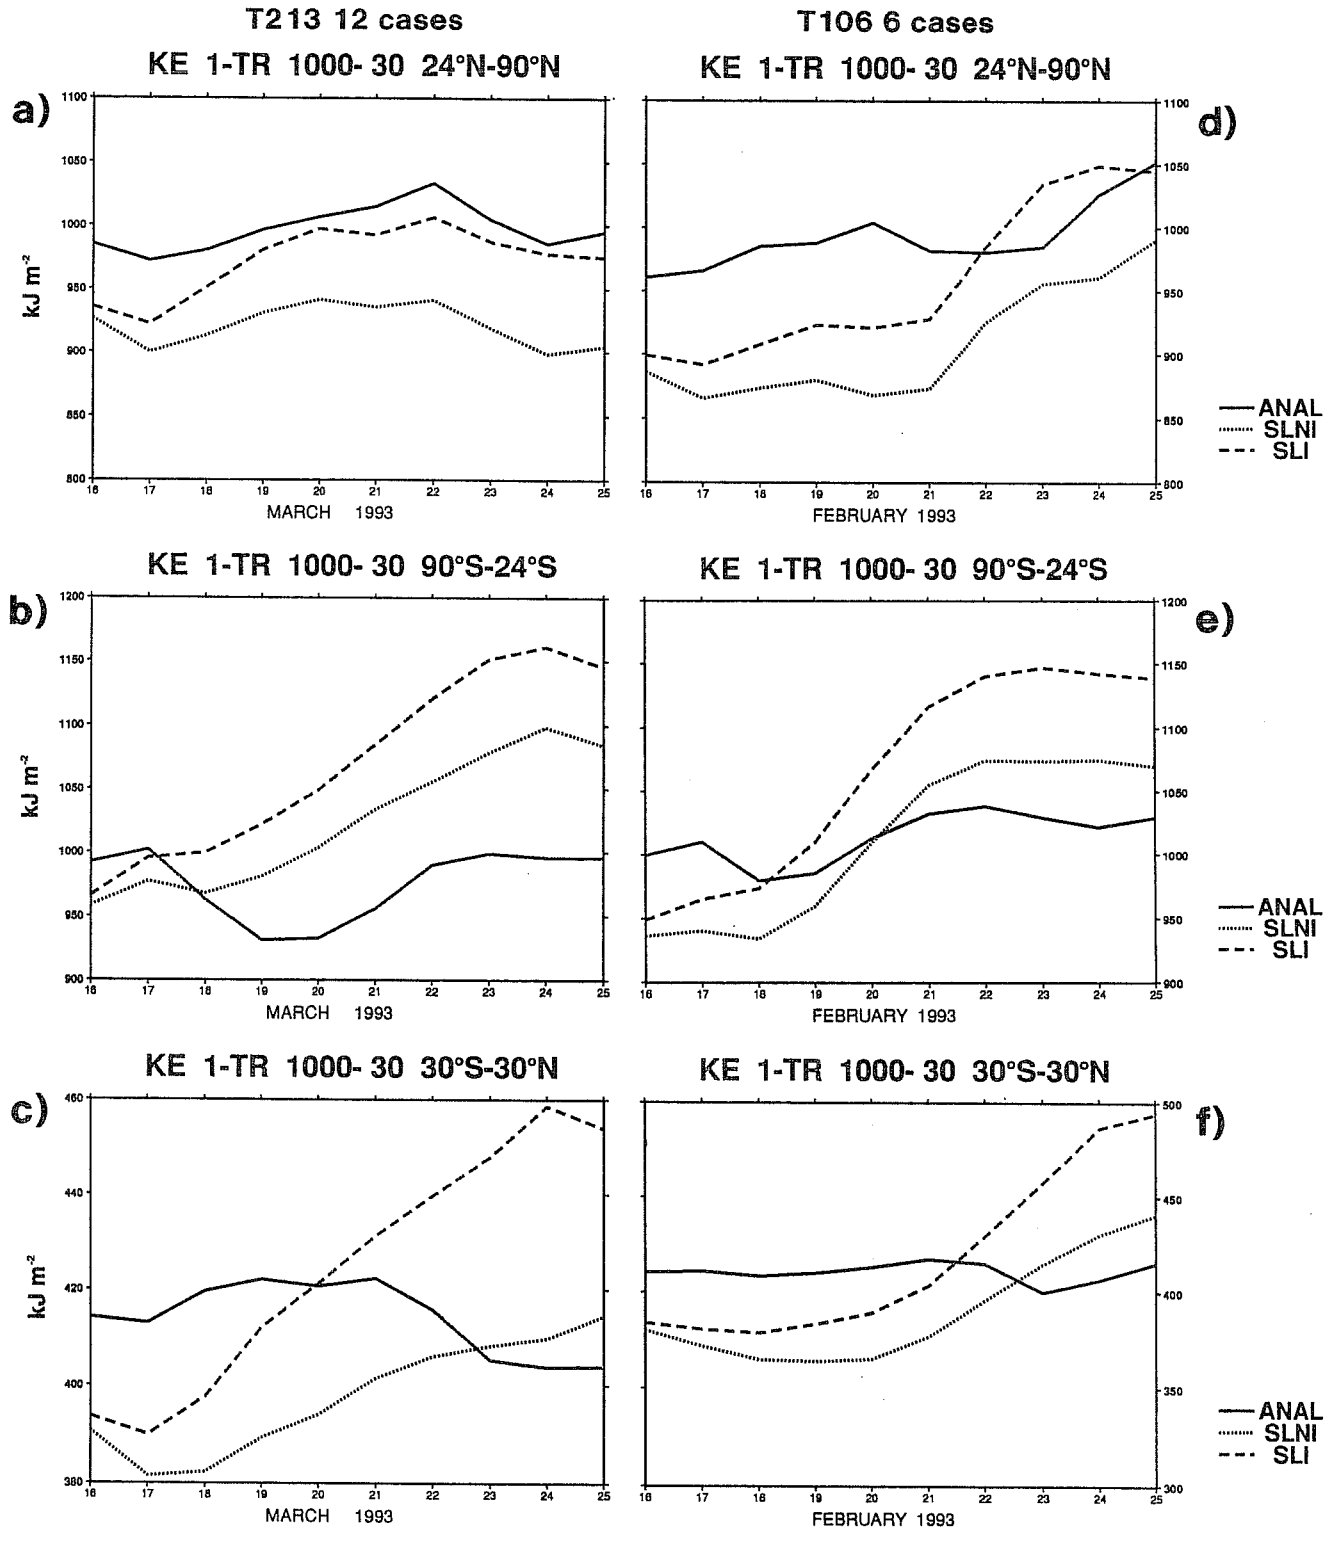

Fig. 3: Mean eddy kinetic energy over three different latitude bands for the T213 experiments of Fig. 2 (panels a-c) and for the T106 experiments of Fig. 1 (panels d-f).

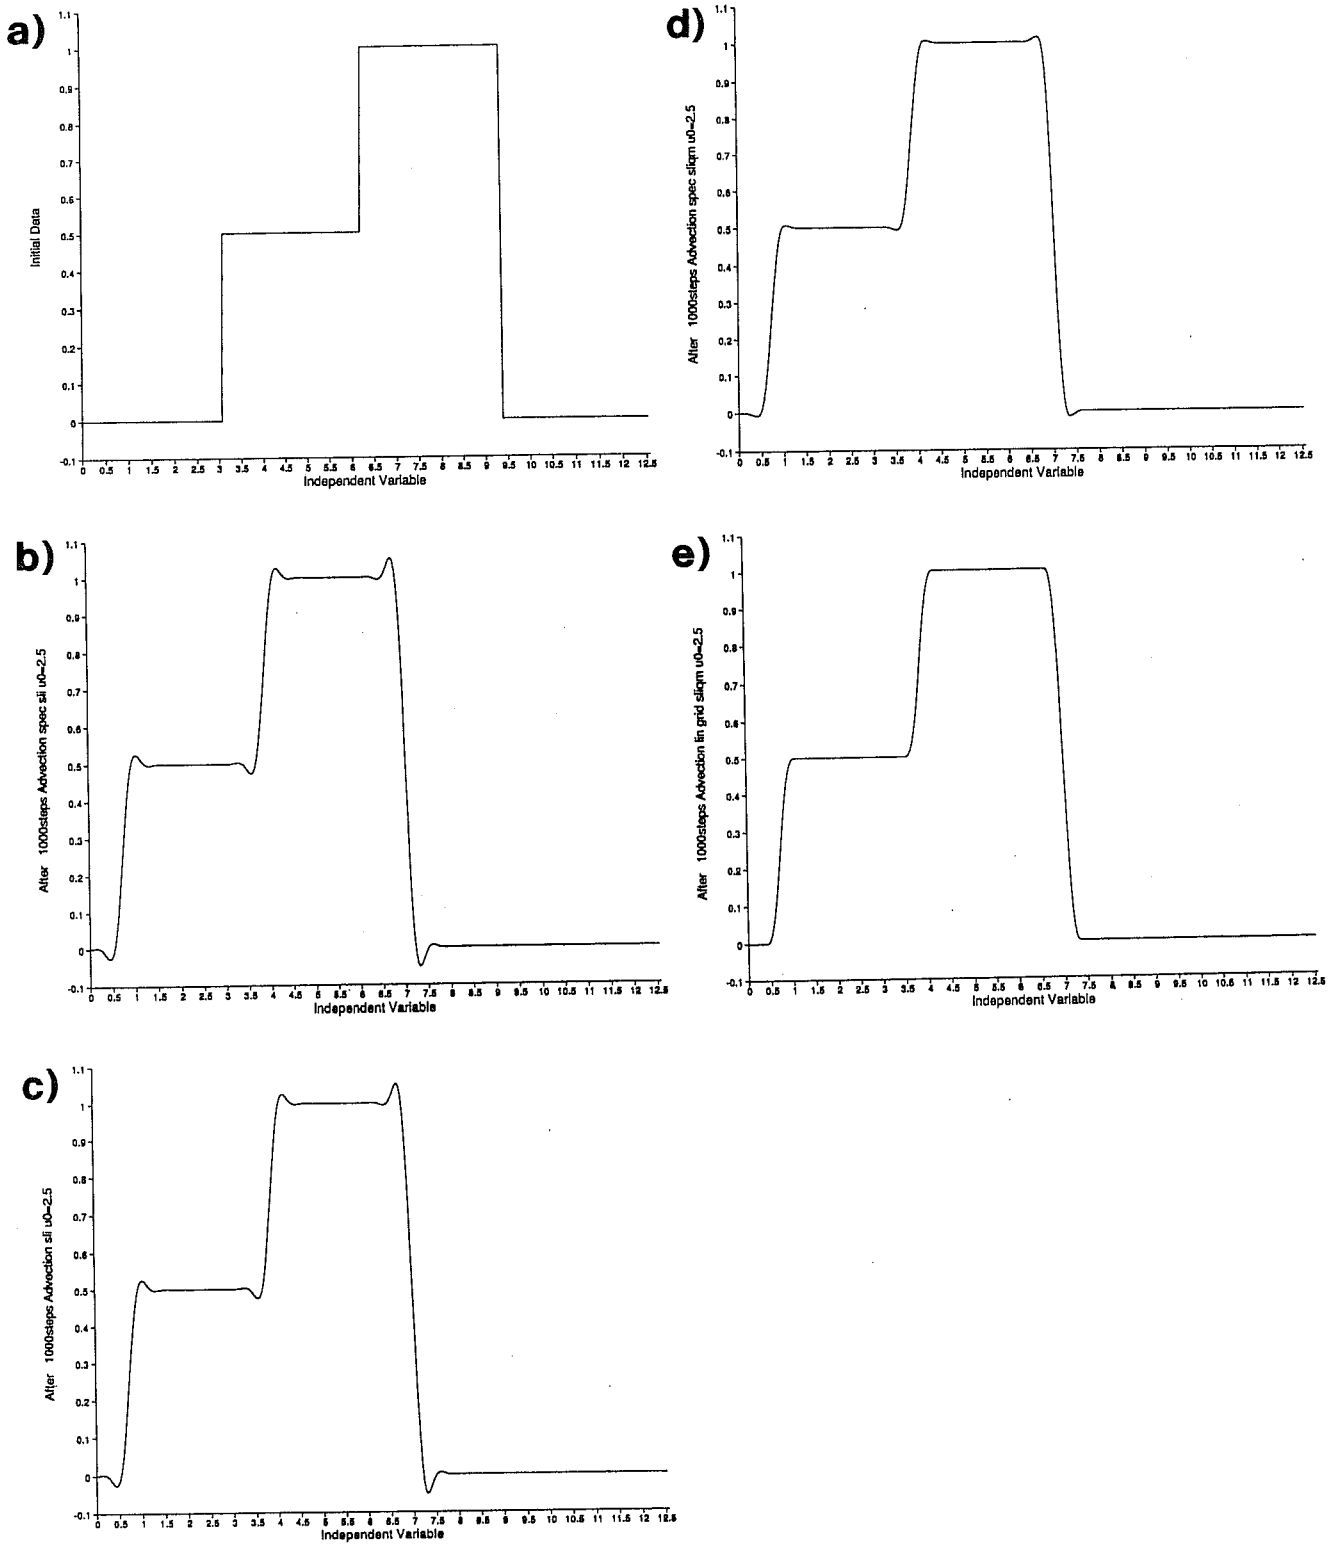

Fig. 4 One-dimensional tests for the linear advection of a sharp field.
(a) Initial field,
(b) after 1000 advection steps using cubic Lagrange interpolation to the departure point,
(c) performing a spectral fitting of the field after each advection step,
(d) using quasi-monotone interpolation to the departure point and spectral fitting,
(e) using quasi-monotone interpolation and the spectral fitting corresponding to the "linear grid".

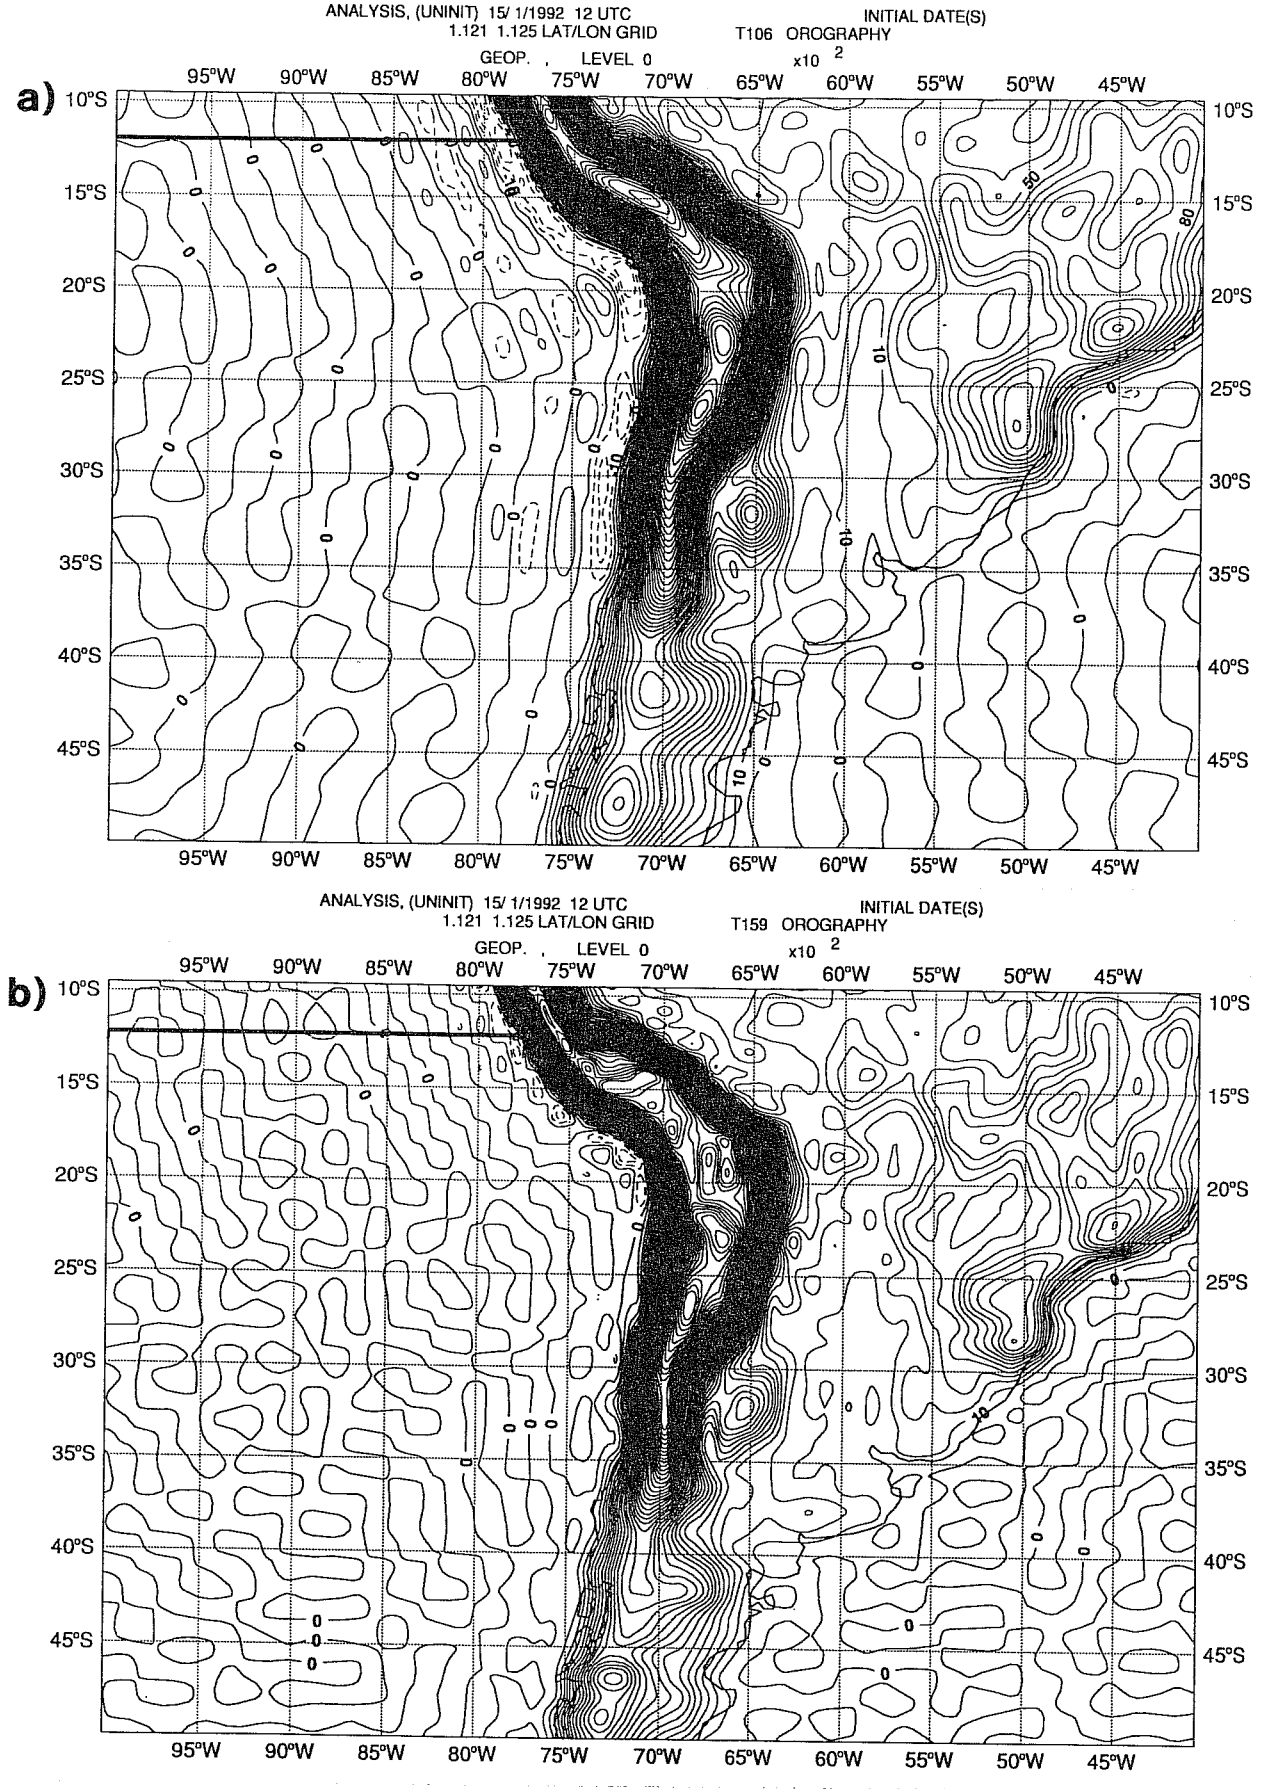

Fig. 5: (a) Spectral T106 orography near the Andes,
(b) spectral T159 orography near the Andes using the linear grid (the Gaussian grid is the same as for panel (a)).

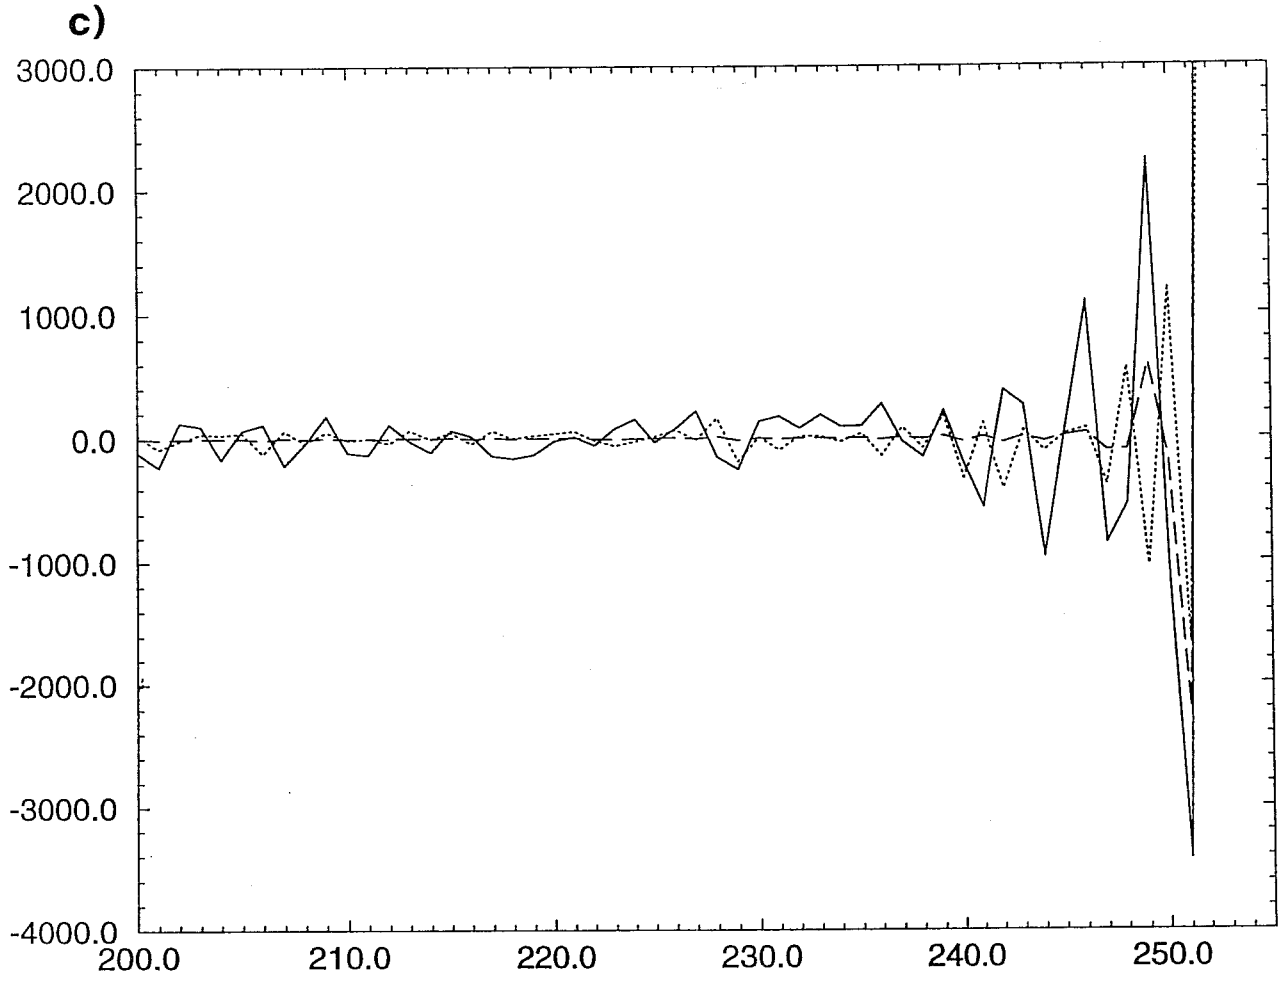

Fig. 5(c): Cross-section of the surface geopotential west of the Andes at 12°S; full line: T106, dotted line: T159, dashed line: "diffused" T159.

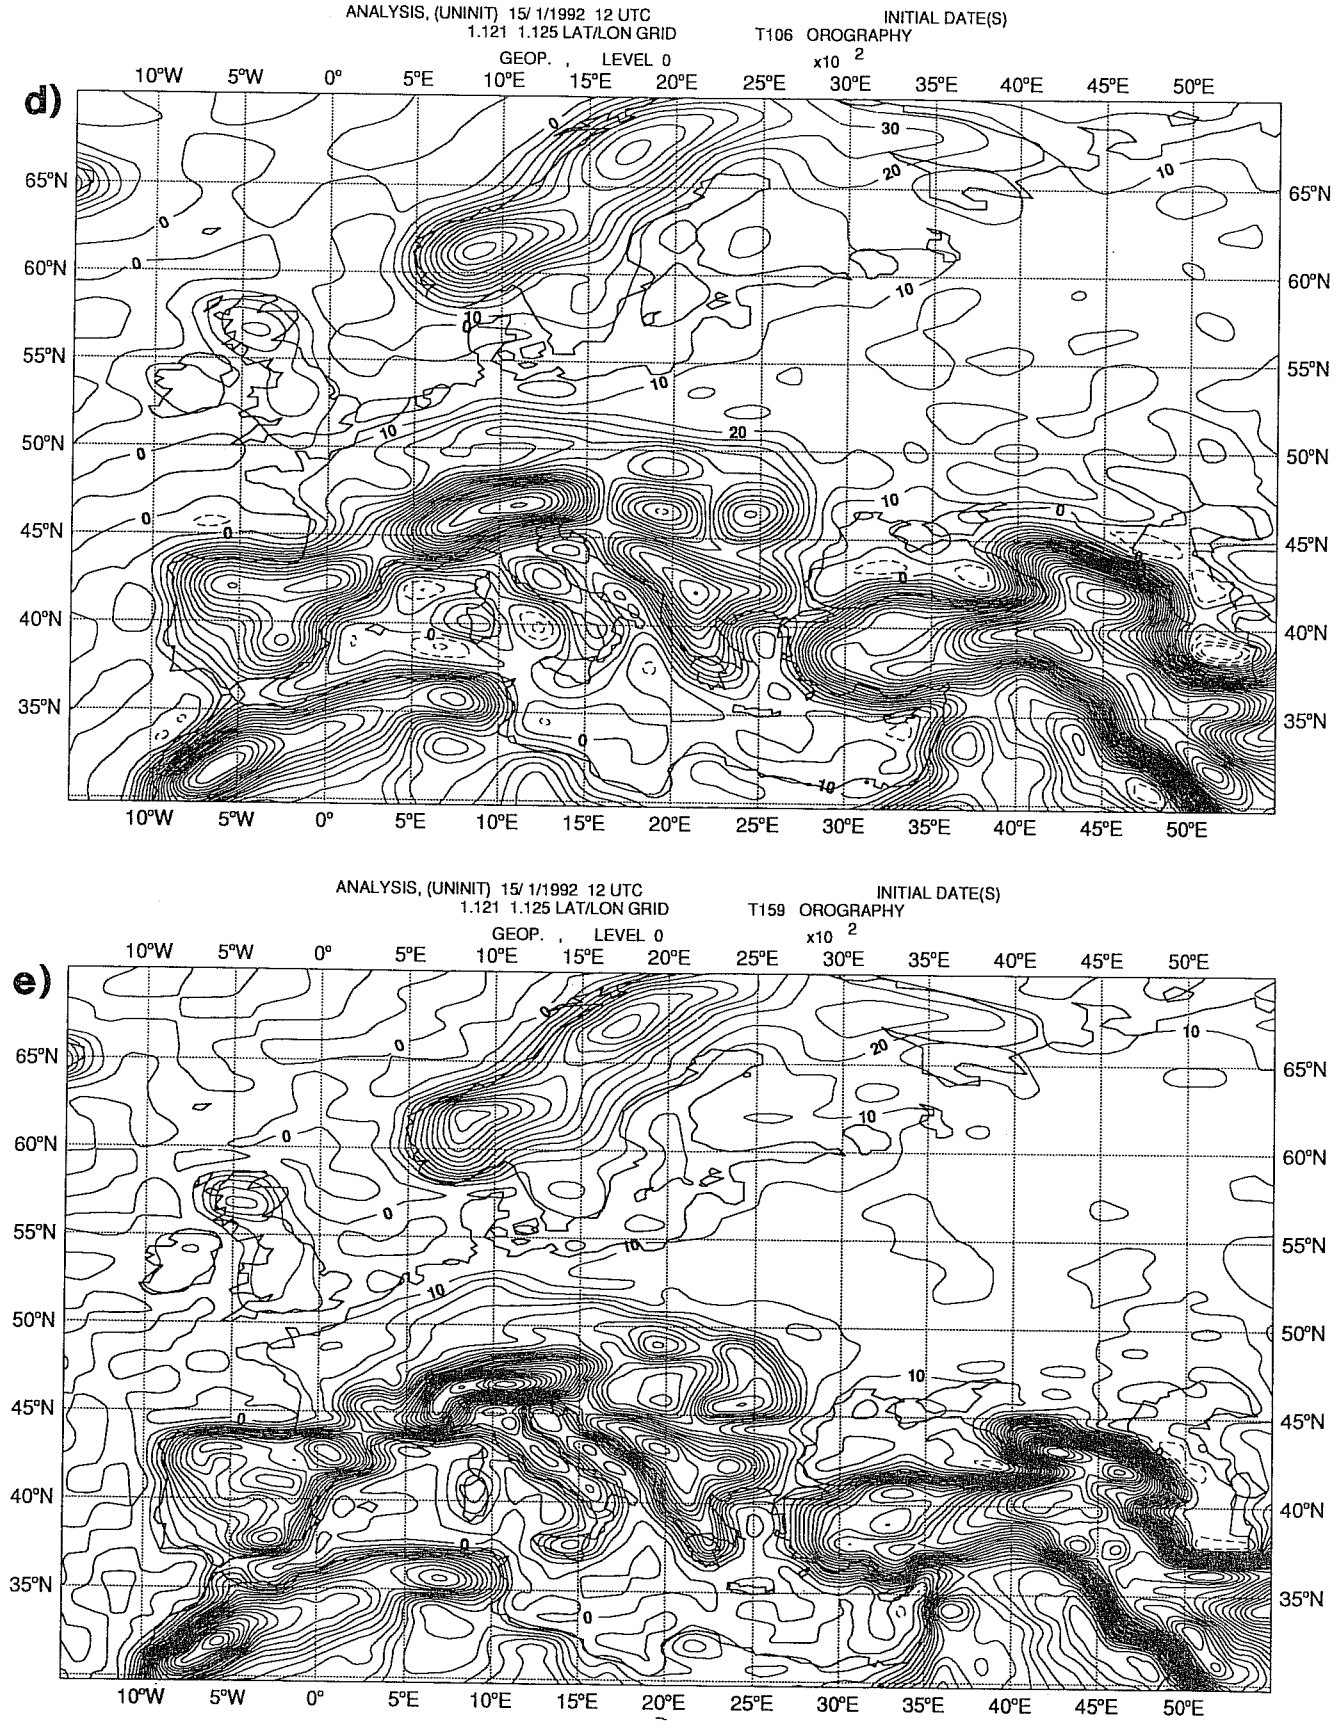

Fig. 5: (d) spectral T106 orography over Europe,
(e) spectral T159 orography over Europe using the linear Gaussian grid.

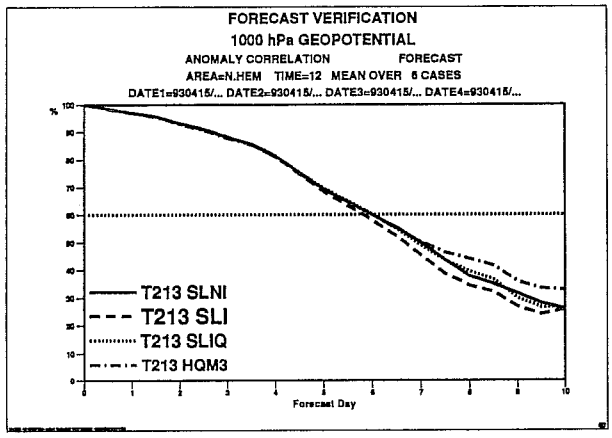

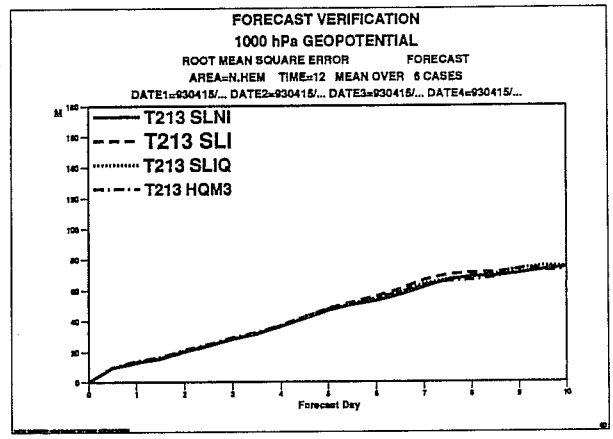

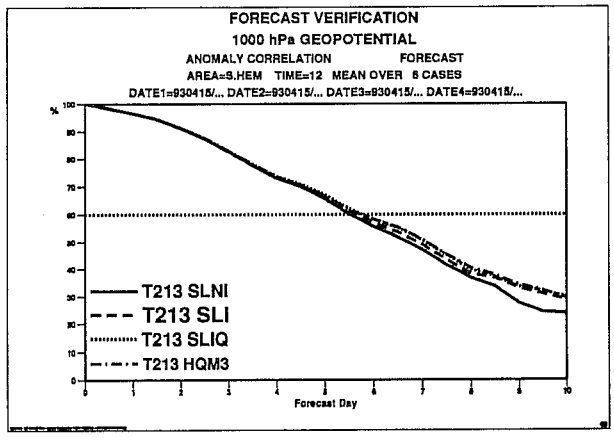

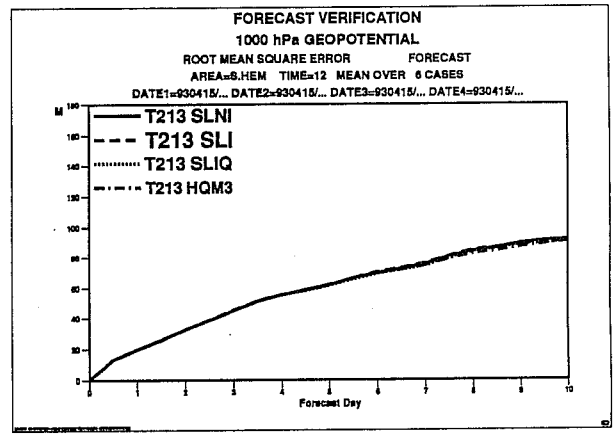

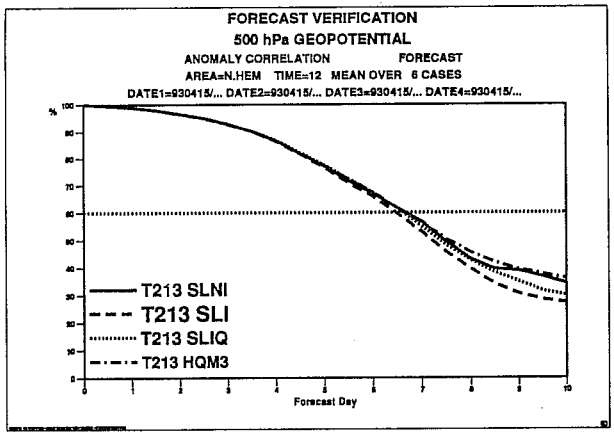

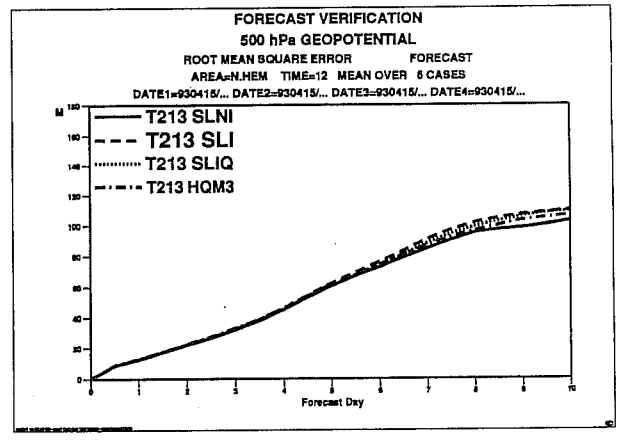

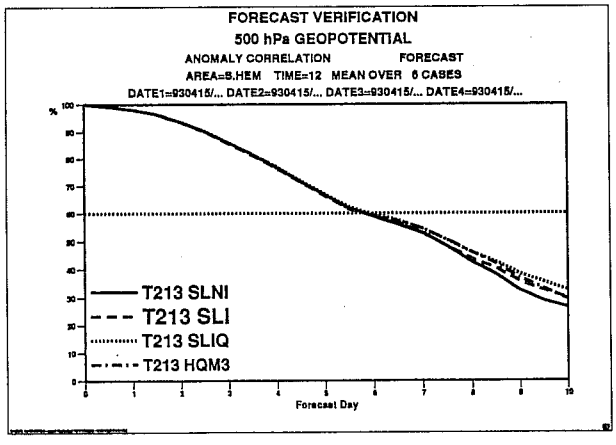

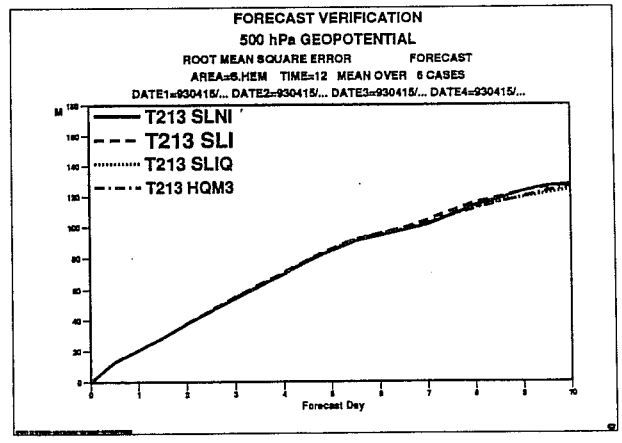

Fig. 6(a): Mean scores of four series of 6 experiments at T213 resolution. SLNI is the standard non-interpolating in the vertical semi-Lagrangian scheme. SLI is the standard fully interpolating semi-Lagrangian scheme. SLIQ is the interpolating semi-Lagrangian scheme using quasi-monotone interpolations in the horizontal and HQM3 is the same as SLIQ but using trilinear interpolation and averaging along the trajectory for the right hand side terms of all the equations.

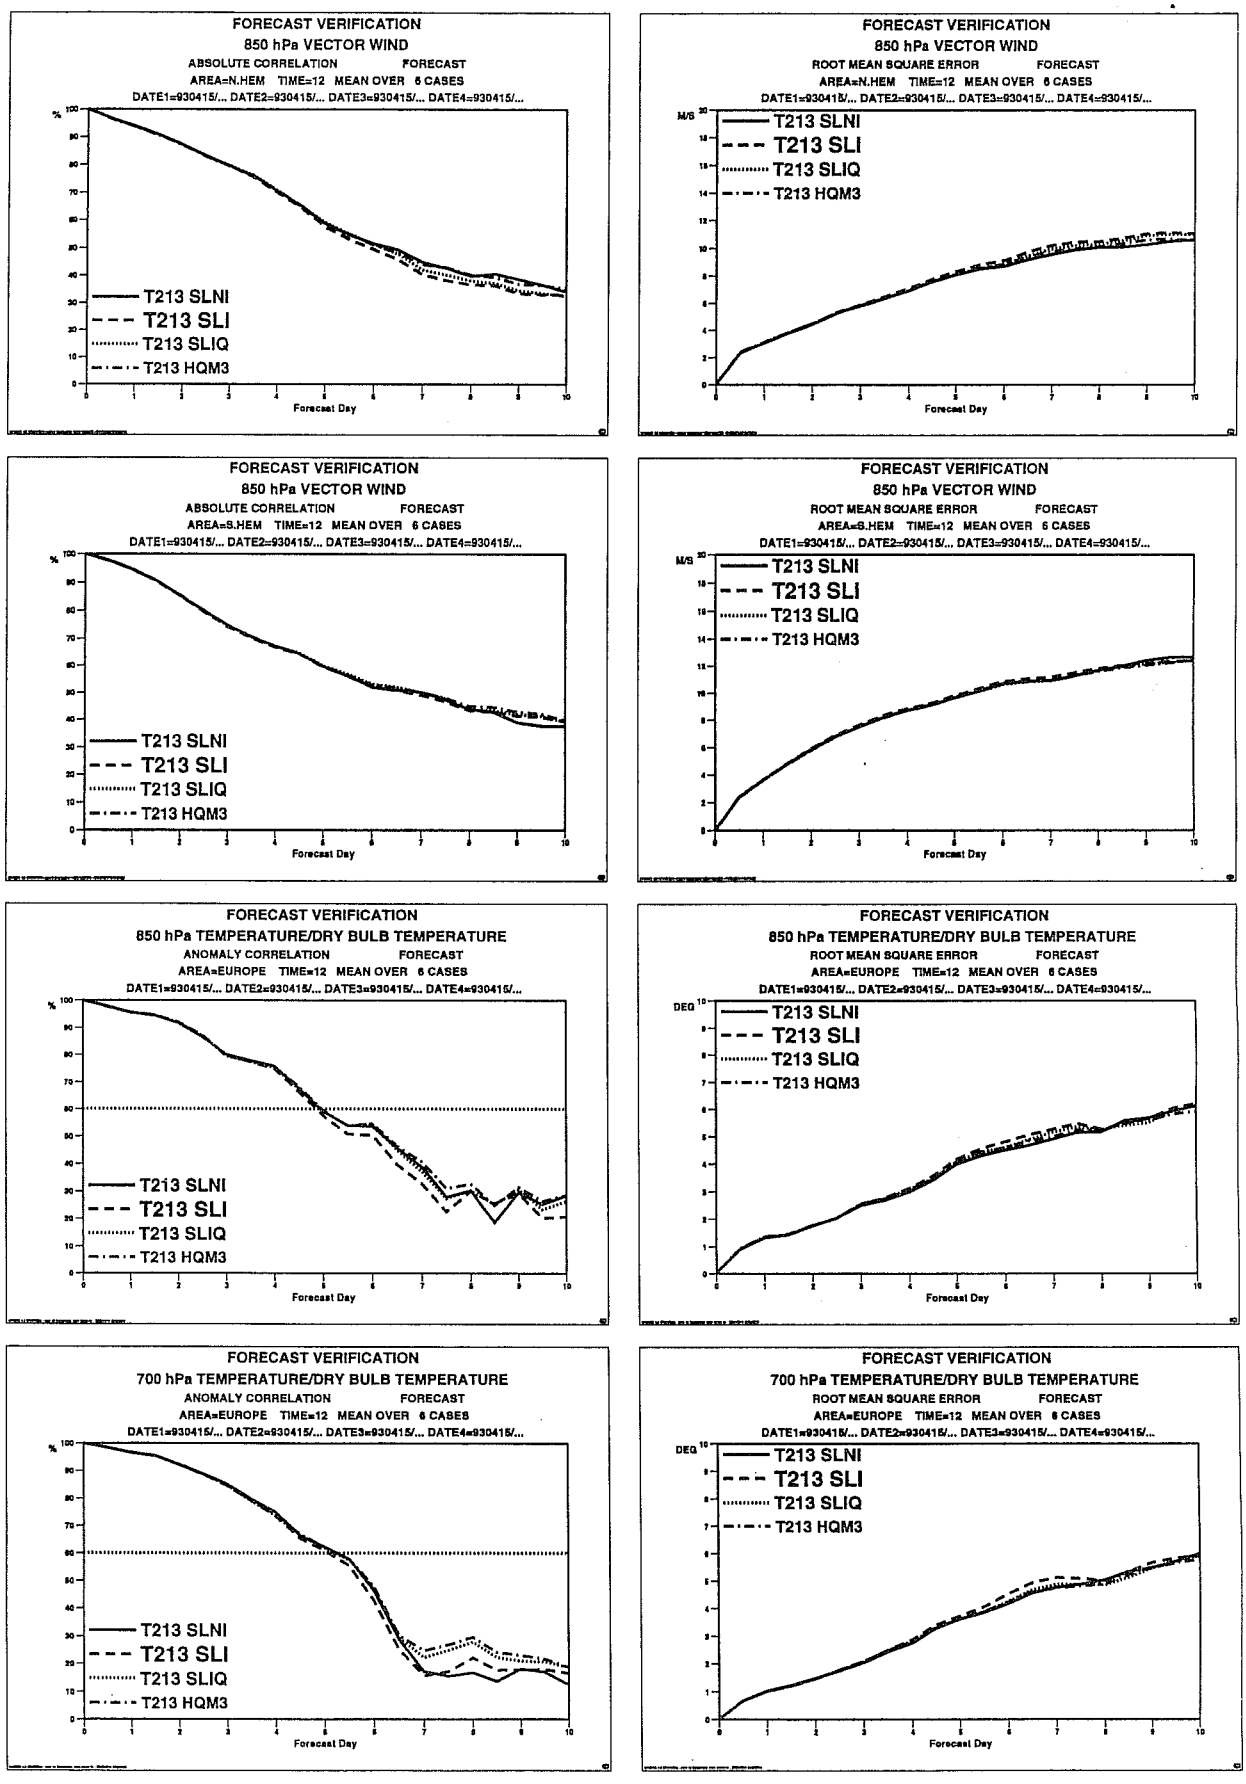

Fig. 6(b): Mean scores of four series of 6 experiments at T213 resolution. SLNI is the standard non-interpolating in the vertical semi-Lagrangian scheme. SLI is the standard fully interpolating semi-Lagrangian scheme. SLIQ is the interpolating semi-Lagrangian scheme using quasi-monotone interpolations in the horizontal and HQM3 is the same as SLIQ but using trilinear interpolation and averaging along the trajectory for the right hand side terms of all the equations.

T213 6 cases

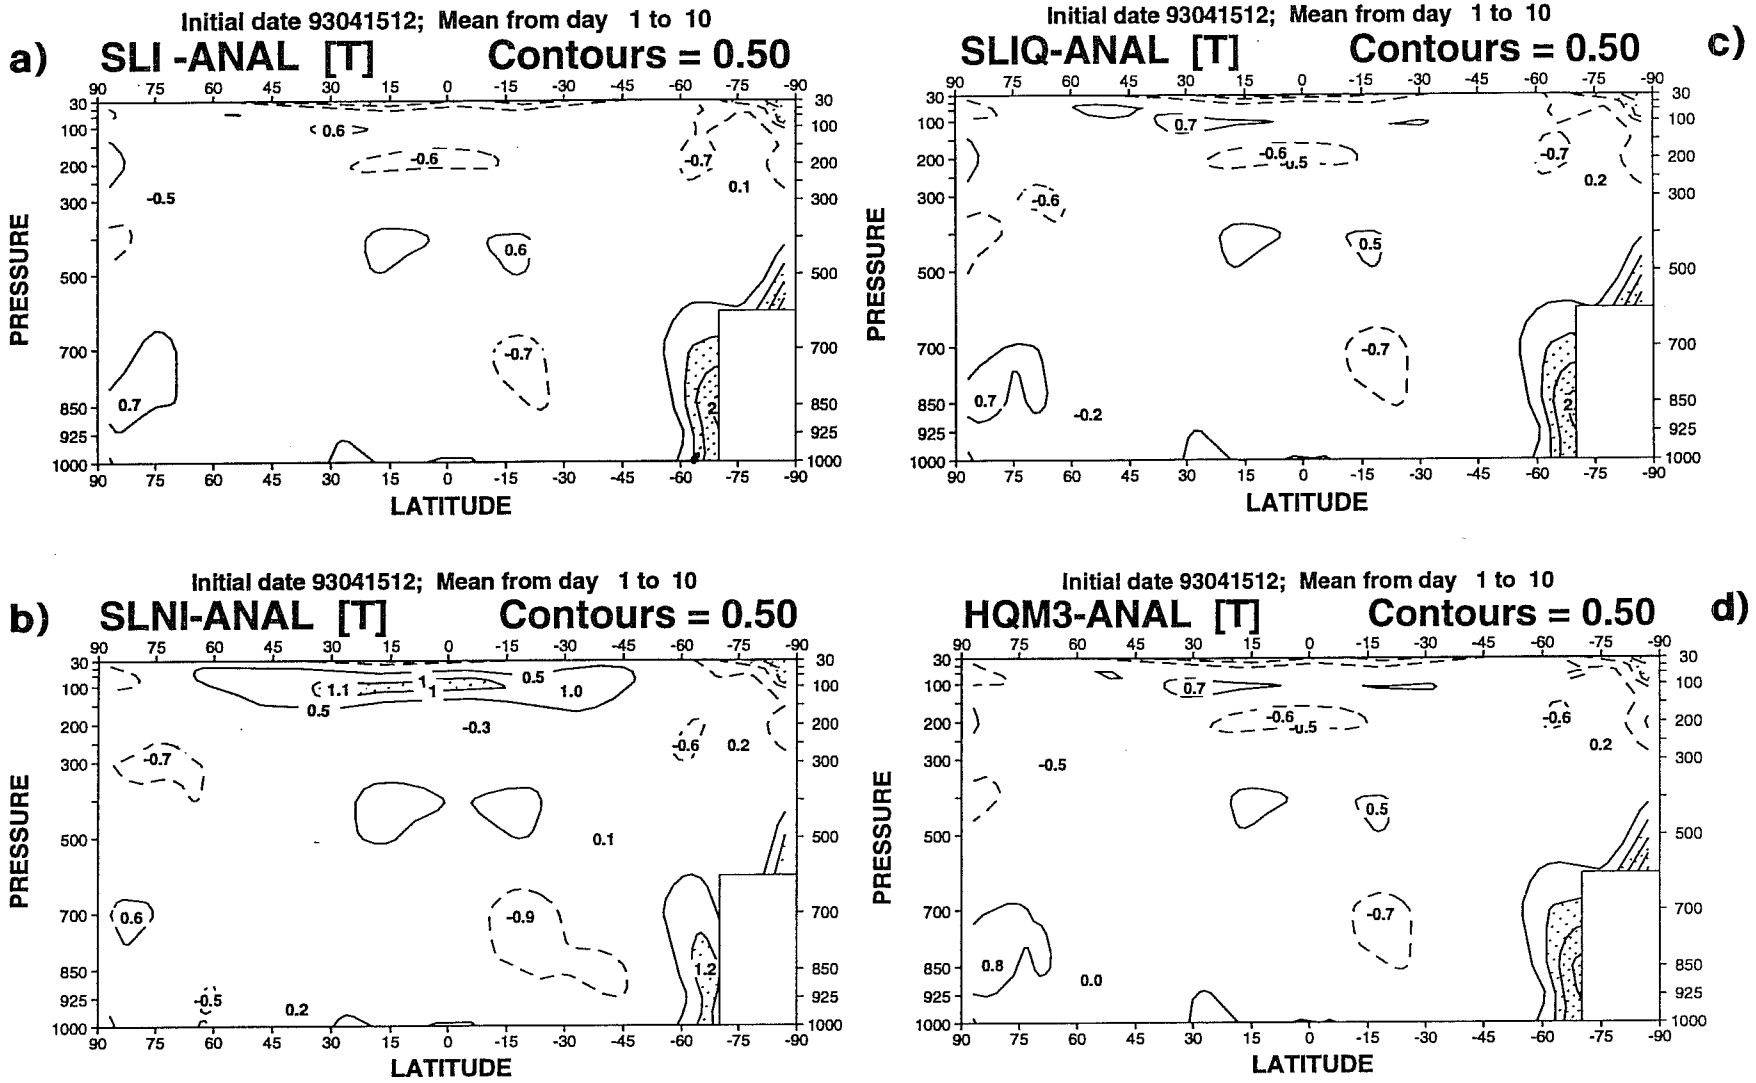

Fig. 7: (a-d) Mean zonal error of the temperature field for the four series of Fig. 6.

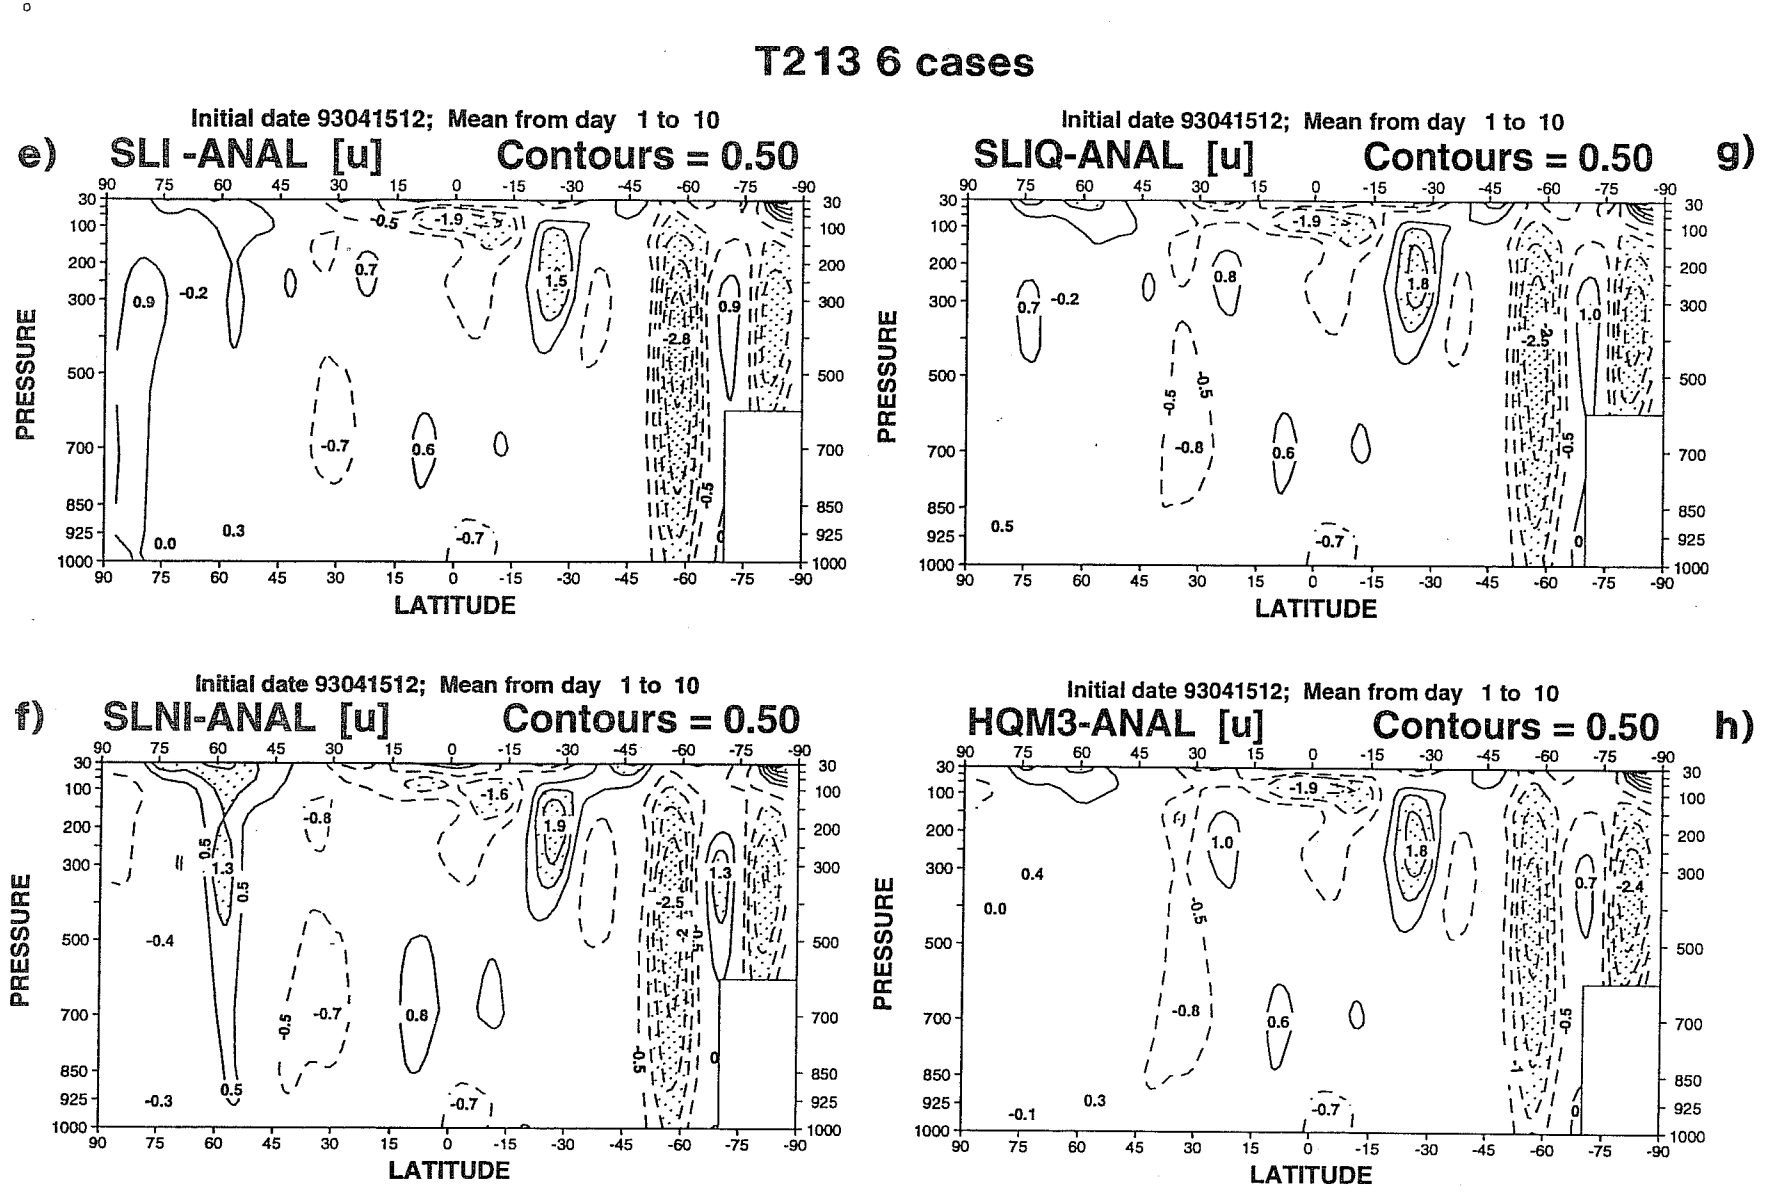

Fig. 7: (e-h) Mean zonal error of the zonal wind field for the four series of Fig. 6.

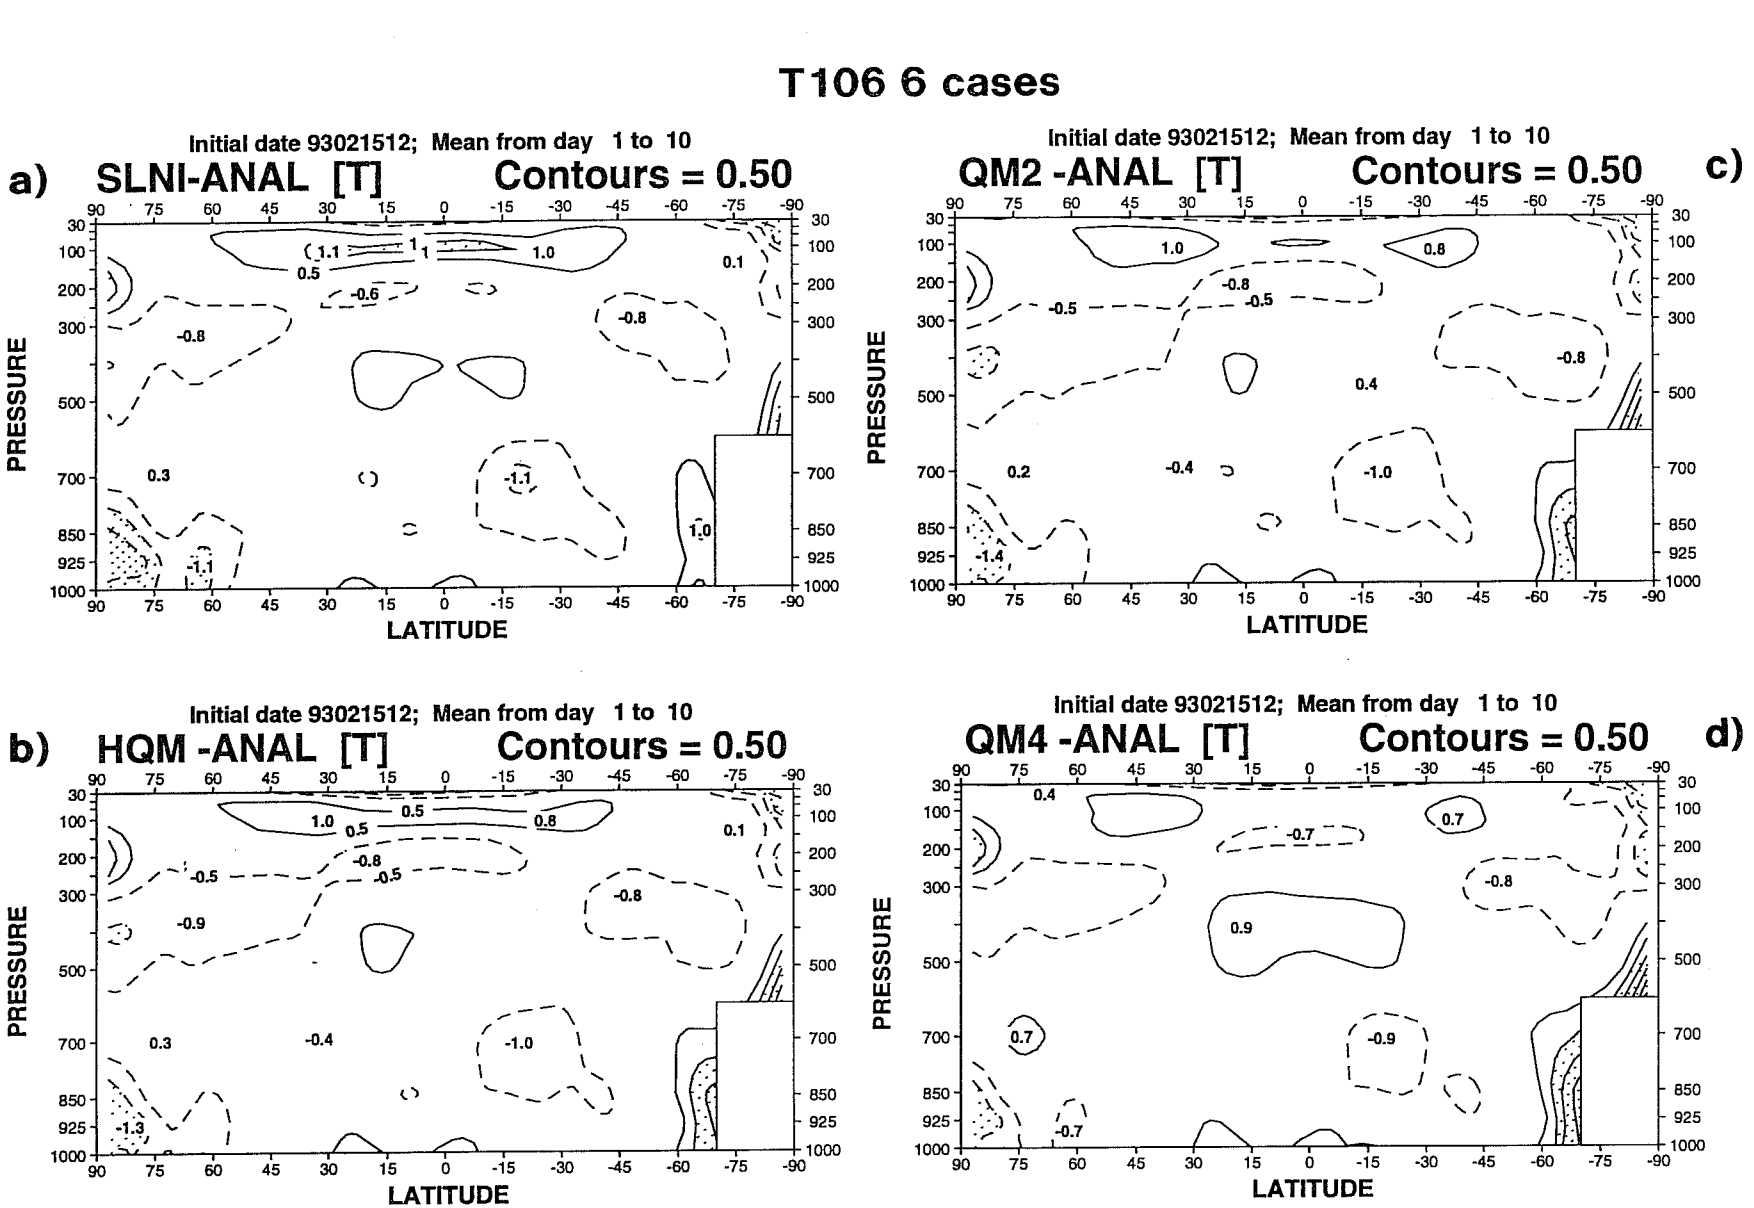

Fig. 8: Mean zonal error of the temperature field for 4 series of T106 experiments.
(a) SLNI: non-interpolating semi-Lagrangian scheme, (b) HQM: interpolating semi-Lagrangian using quasi-monotone interpolation in the horizontal for all the fields,
(c) QM2: as HQM but quasi-monotone interpolation also in the vertical for the wind fields.
(d) QM4: as QM2 but using also a "sharpening" interpolation in the vertical for the temperature field.

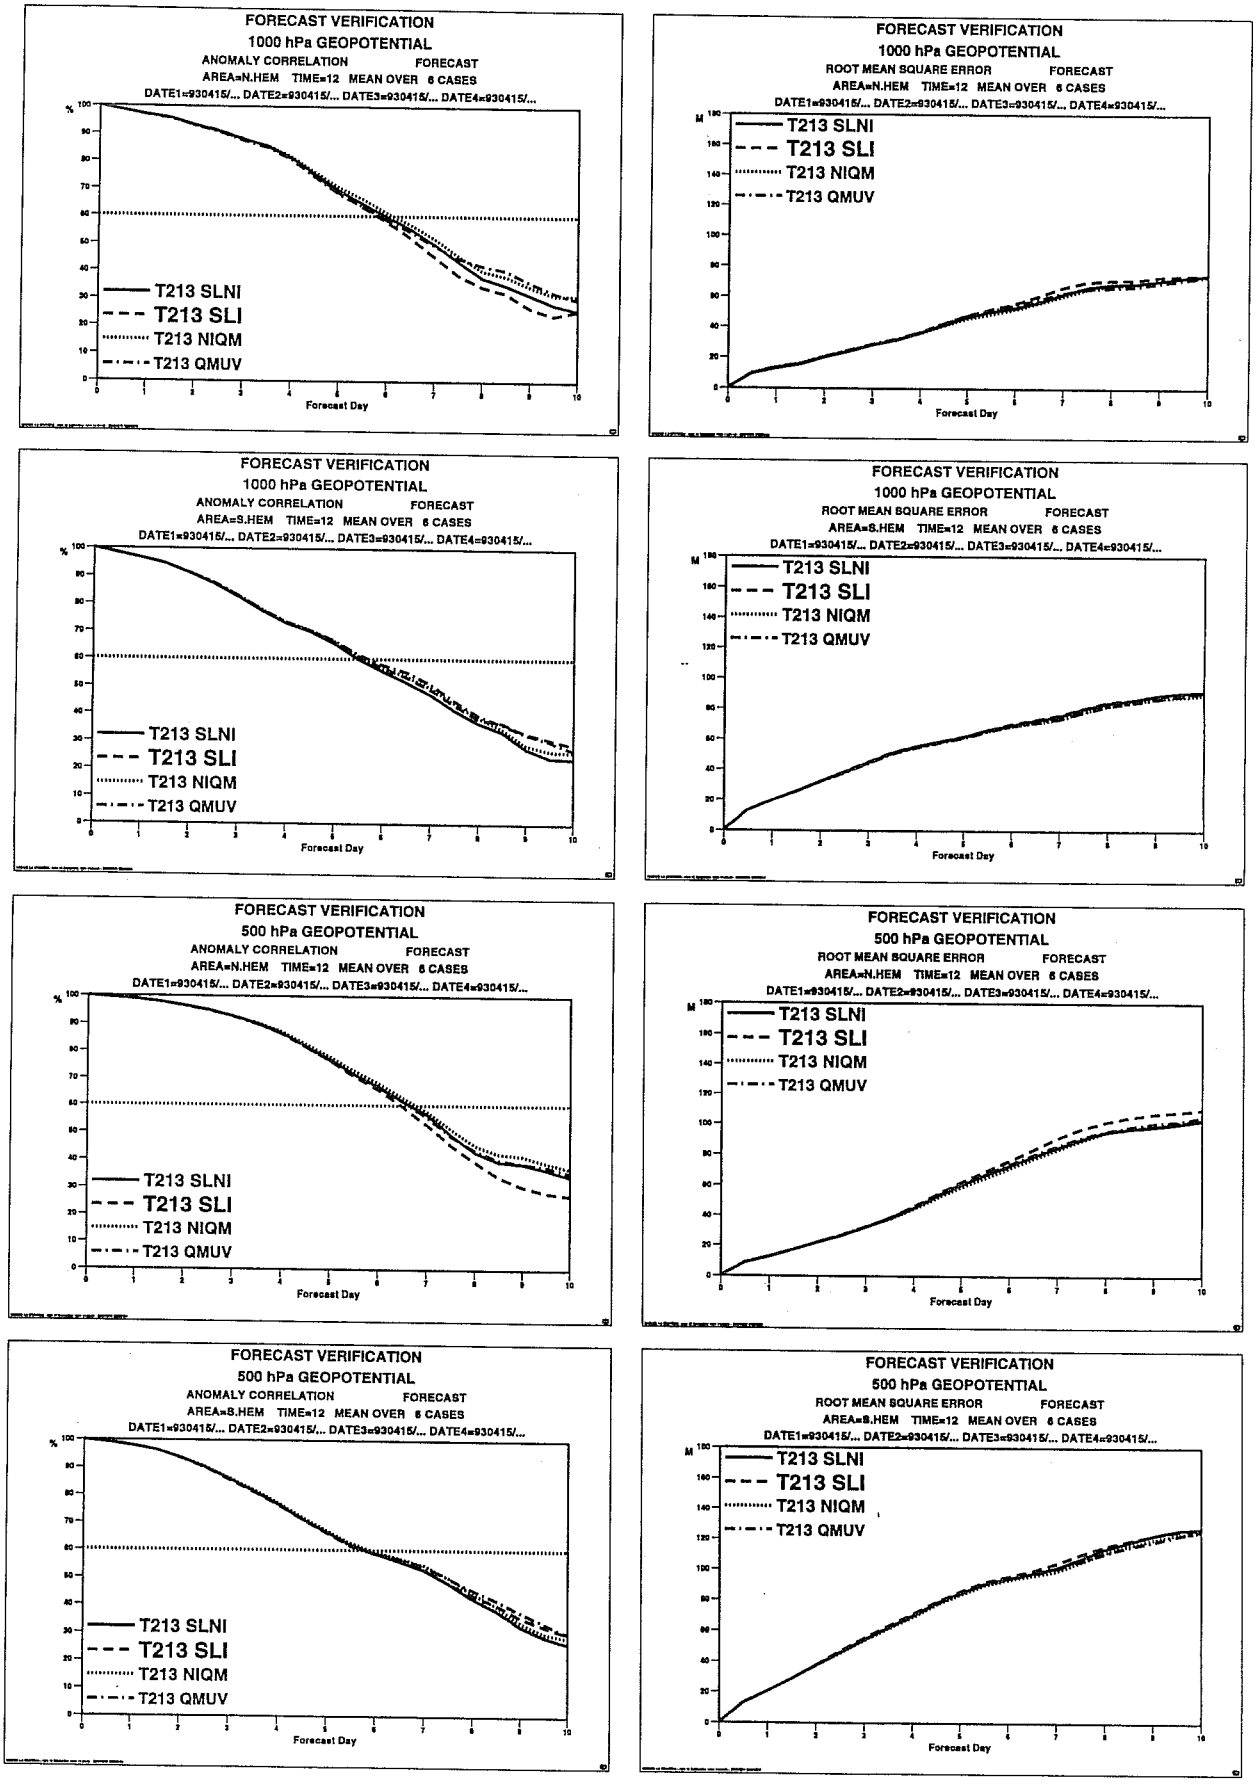

Fig. 9(a): Mean scores for 4 series of experiments at T213 resolution. SLNI is the non-interpolating in the vertical scheme. SLI is the fully interpolating scheme. NIQM is the SLNI scheme using quasi-monotone interpolations. QMUV is the SLI scheme using three-dimensional quasi-monotone interpolations for the wind fields.

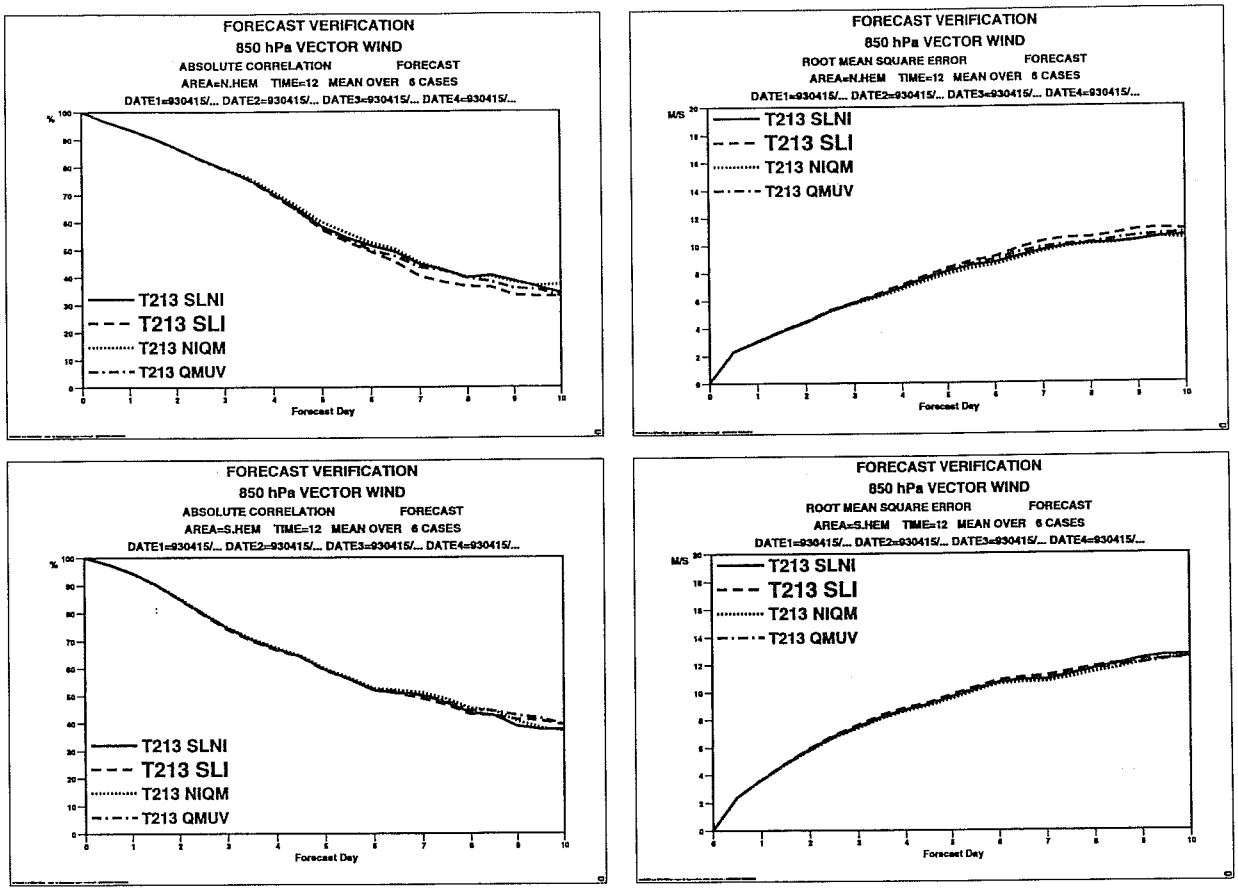

Fig. 9(b): Mean scores for 4 series of experiments at T213 resolution. SLNI is the non-interpolating in the vertical scheme. SLI is the fully interpolating scheme. NIQM is the SLNI scheme using quasi-monotone interpolations. QMUV is the SLI scheme using three-dimensional quasi-monotone interpolations for the wind fields.

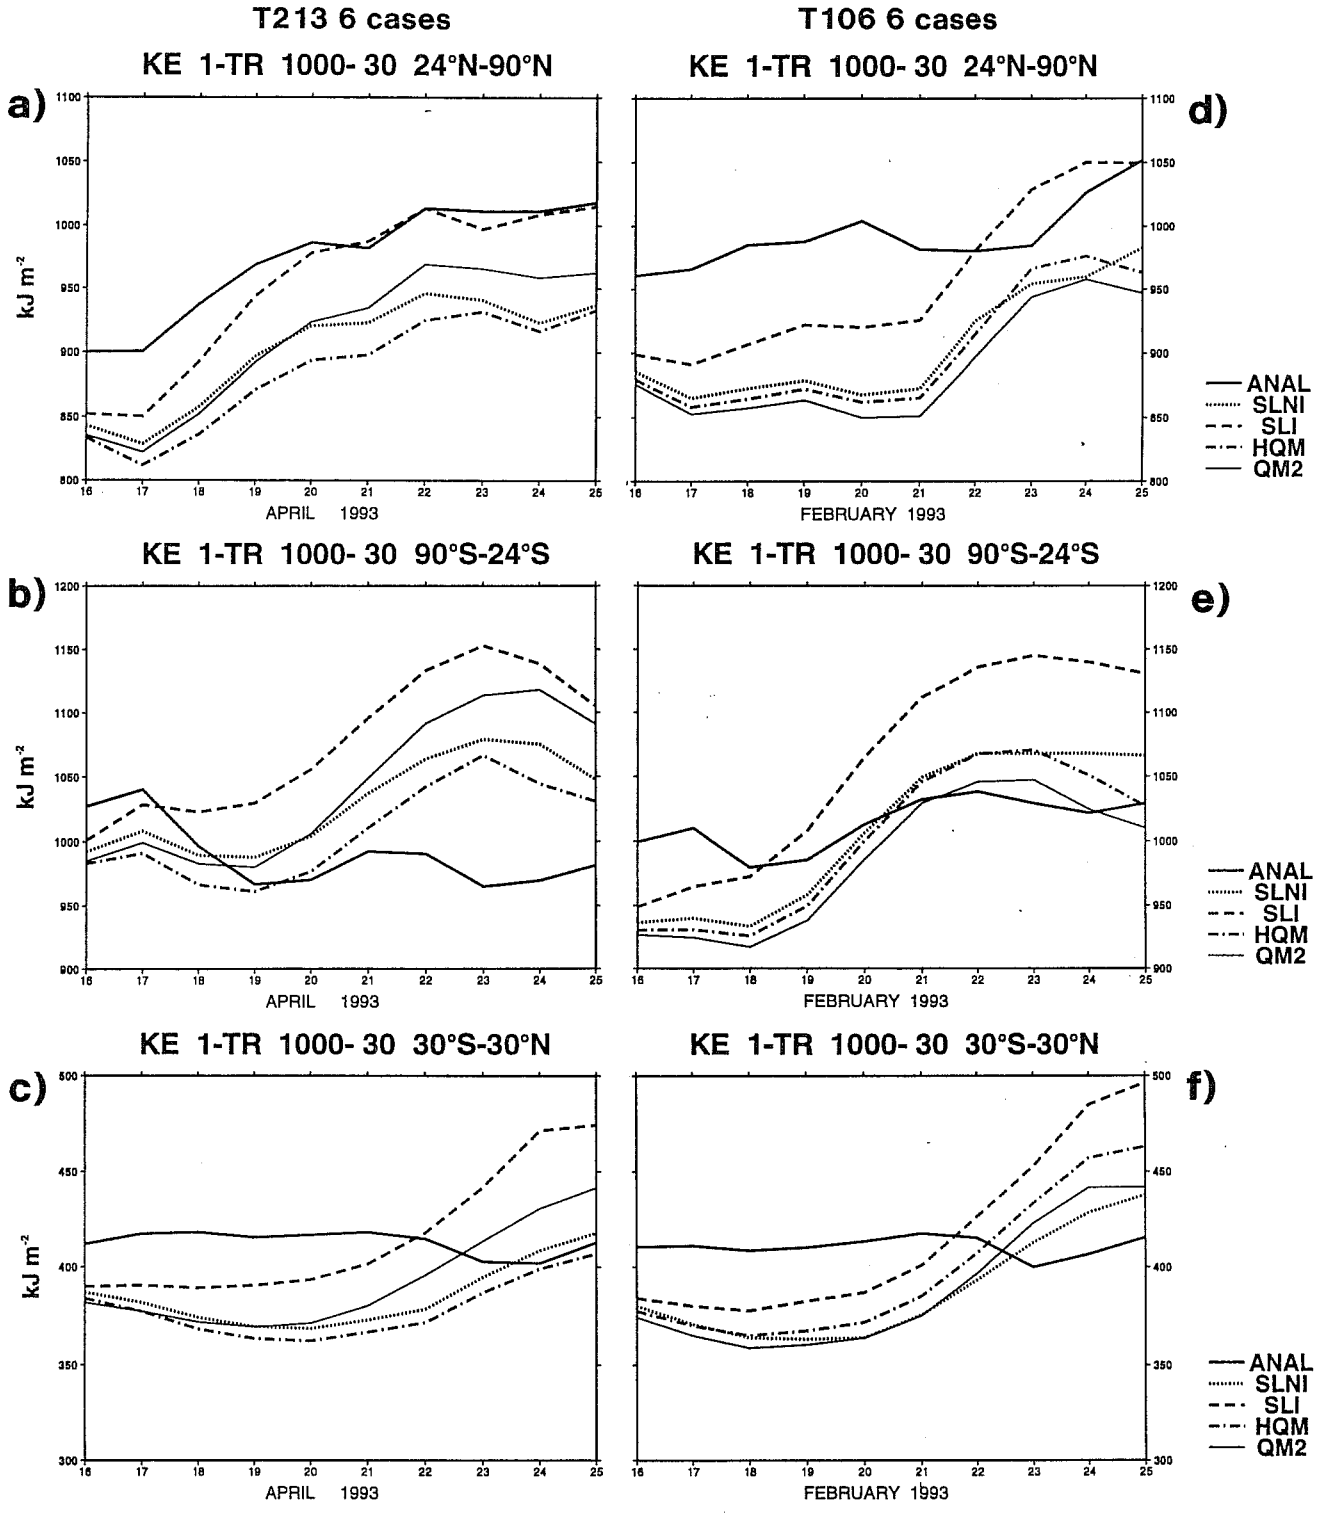

Fig. 10: (a-c) Mean eddy kinetic energy for the experiments of Fig. 9,
(d-f) for the experiments of Fig. 8.

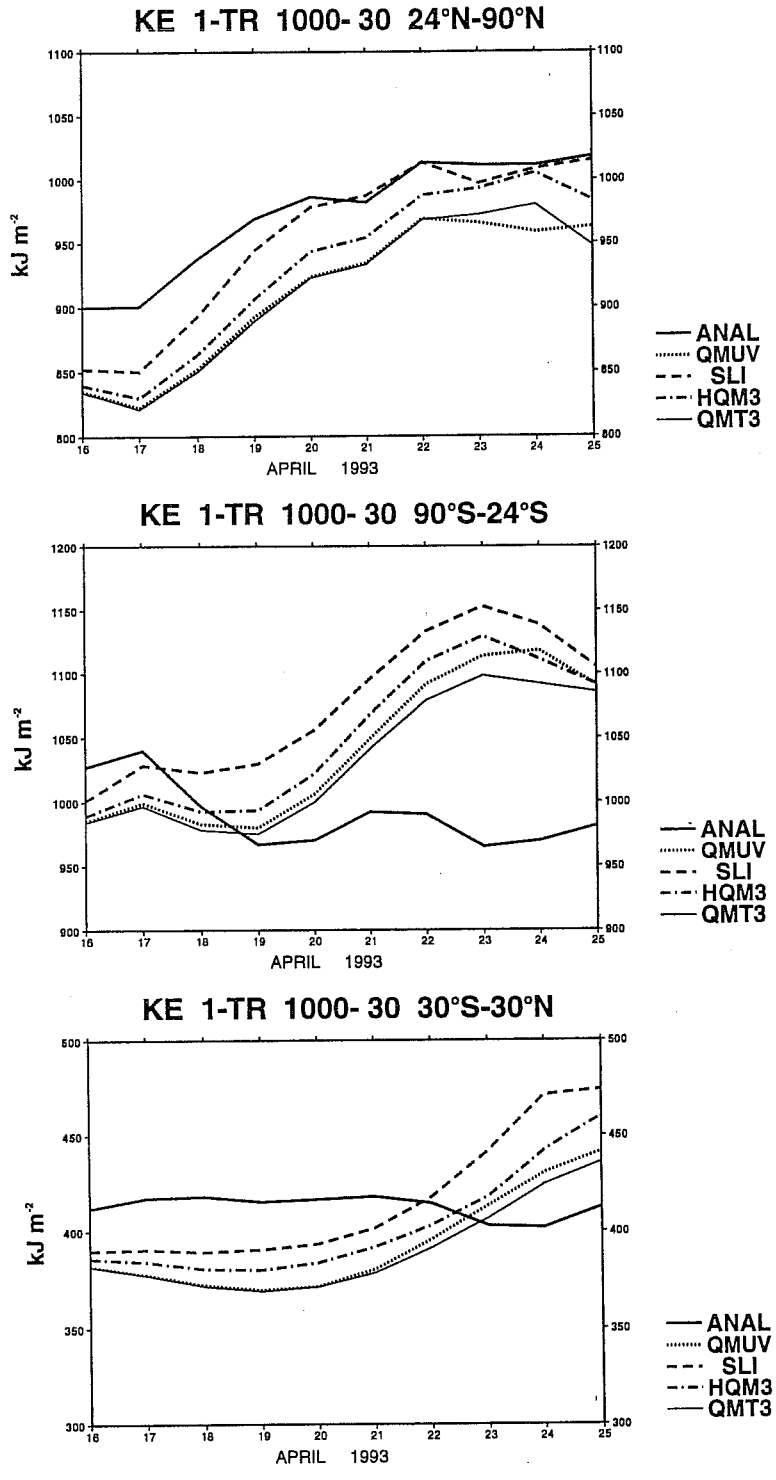

Fig. 11: Mean eddy kinetic energy of four series of experiments at T213 showing the relative importance of the quasi-monotone interpolation applied to the horizontal or to the vertical direction.

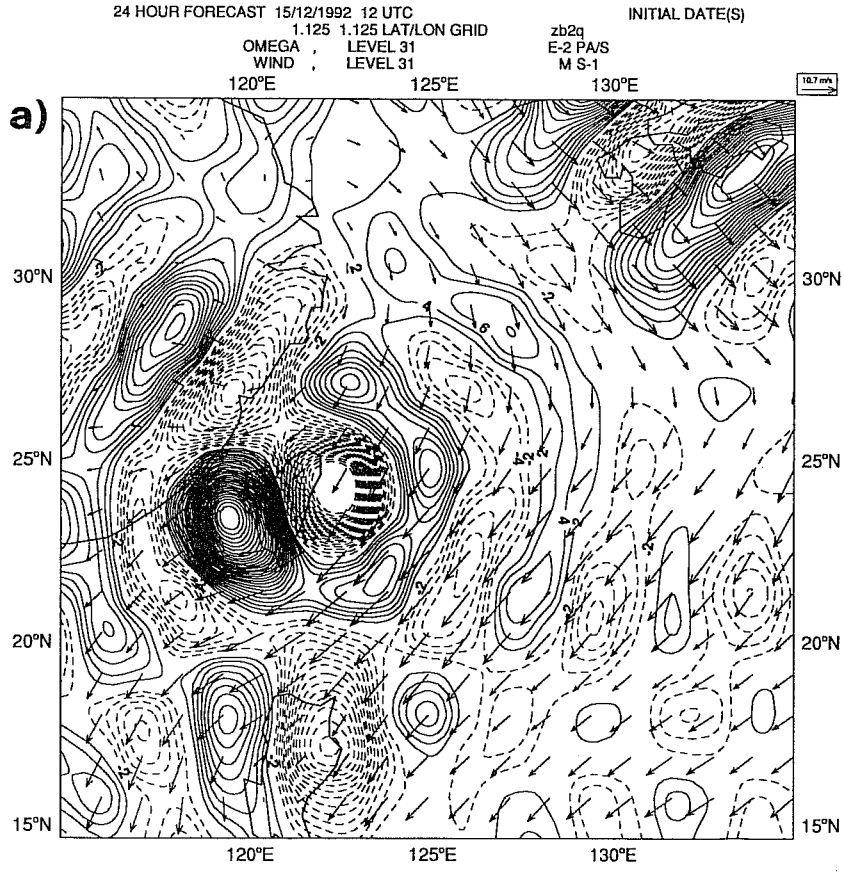

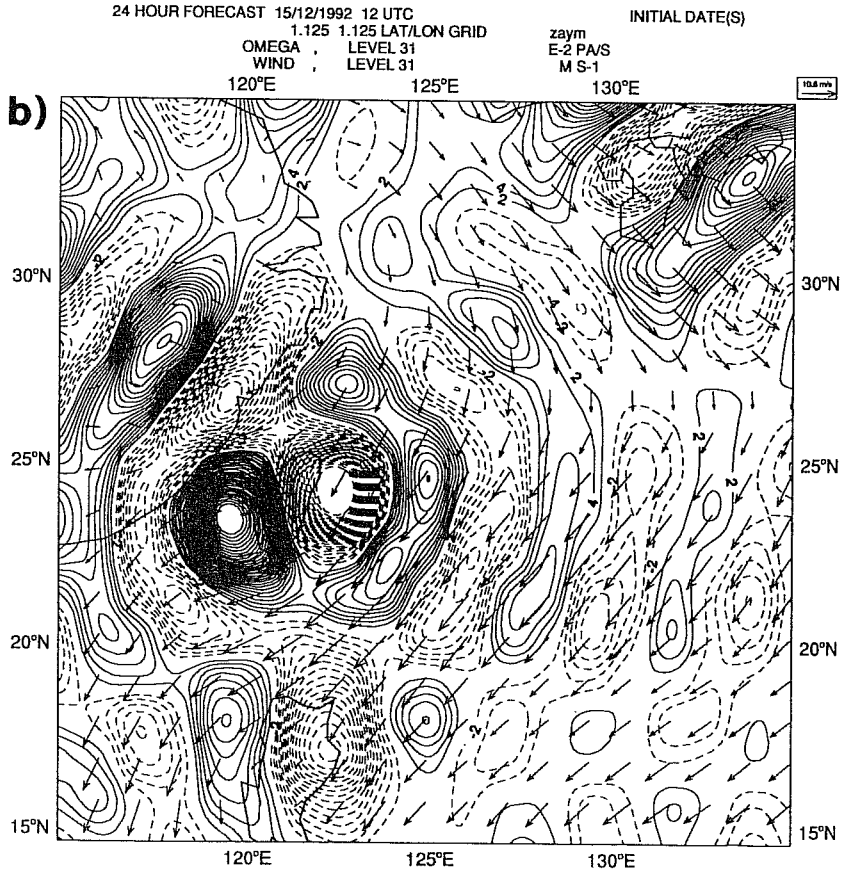

Fig. 12: Vertical velocity and horizontal wind at the lowest model level.
(a) experiment at T106 resolution,
(b) experiment at T159 resolution using the "linear grid" and the same orography as the T106 experiment.

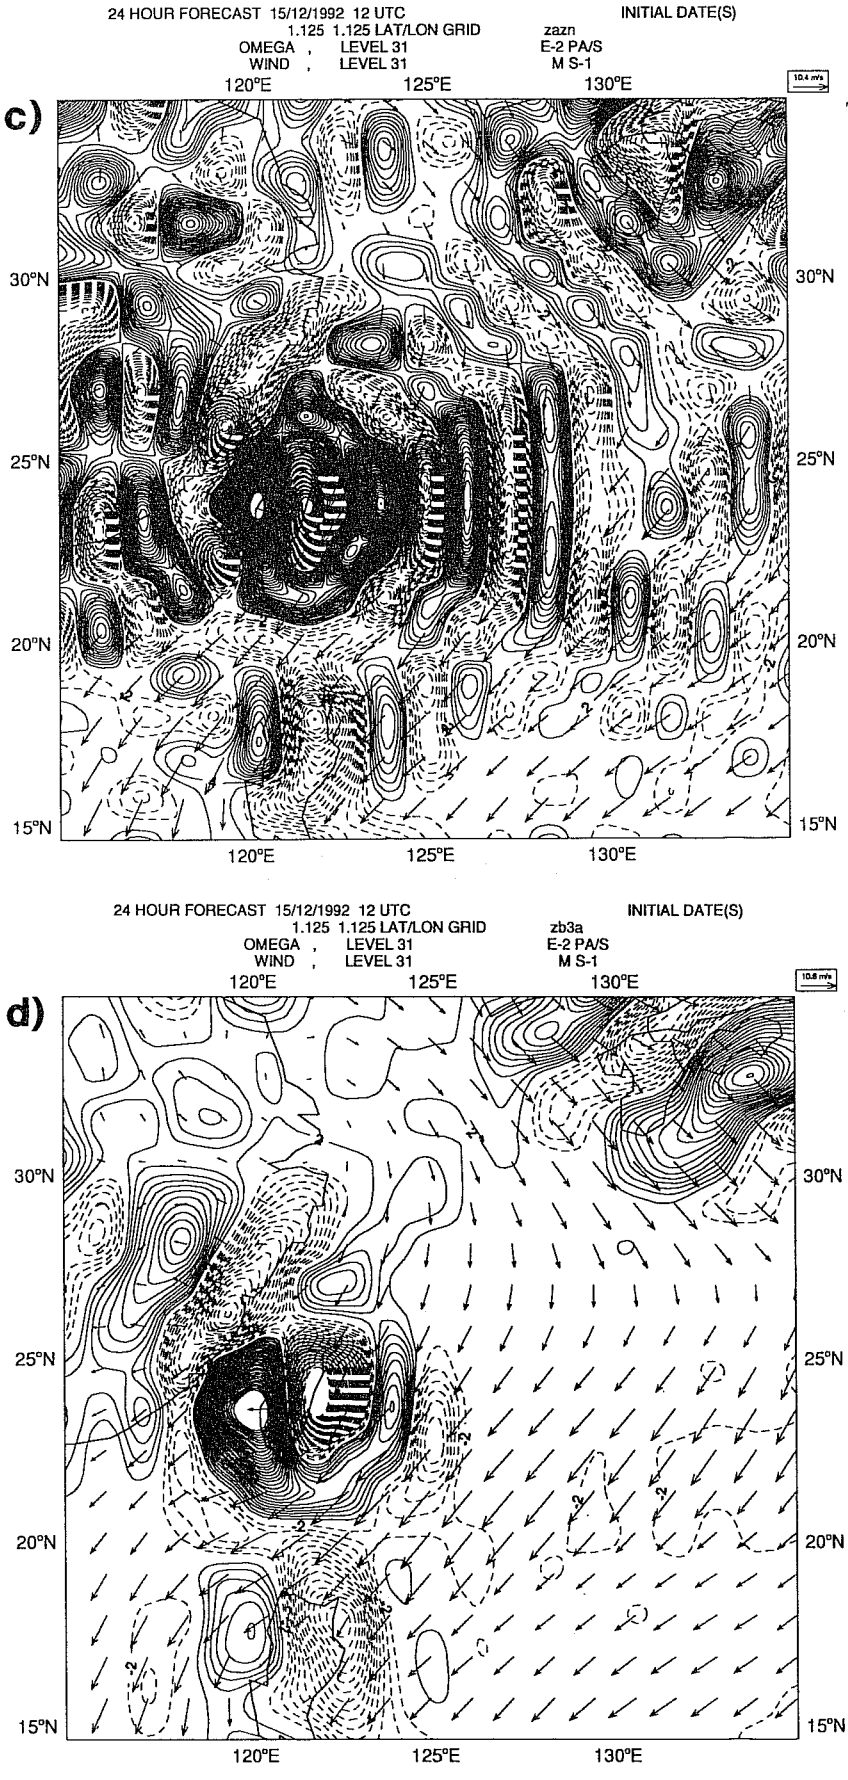

Fig. 12: Vertical velocity and horizontal wind at the lowest model level.
(c) experiment at T159 resolution with the linear grid and orography spectrally fitted at T159 resolution,
(d) experiment at T159 resolution and orography smoothed by means of a high order diffusion operator.

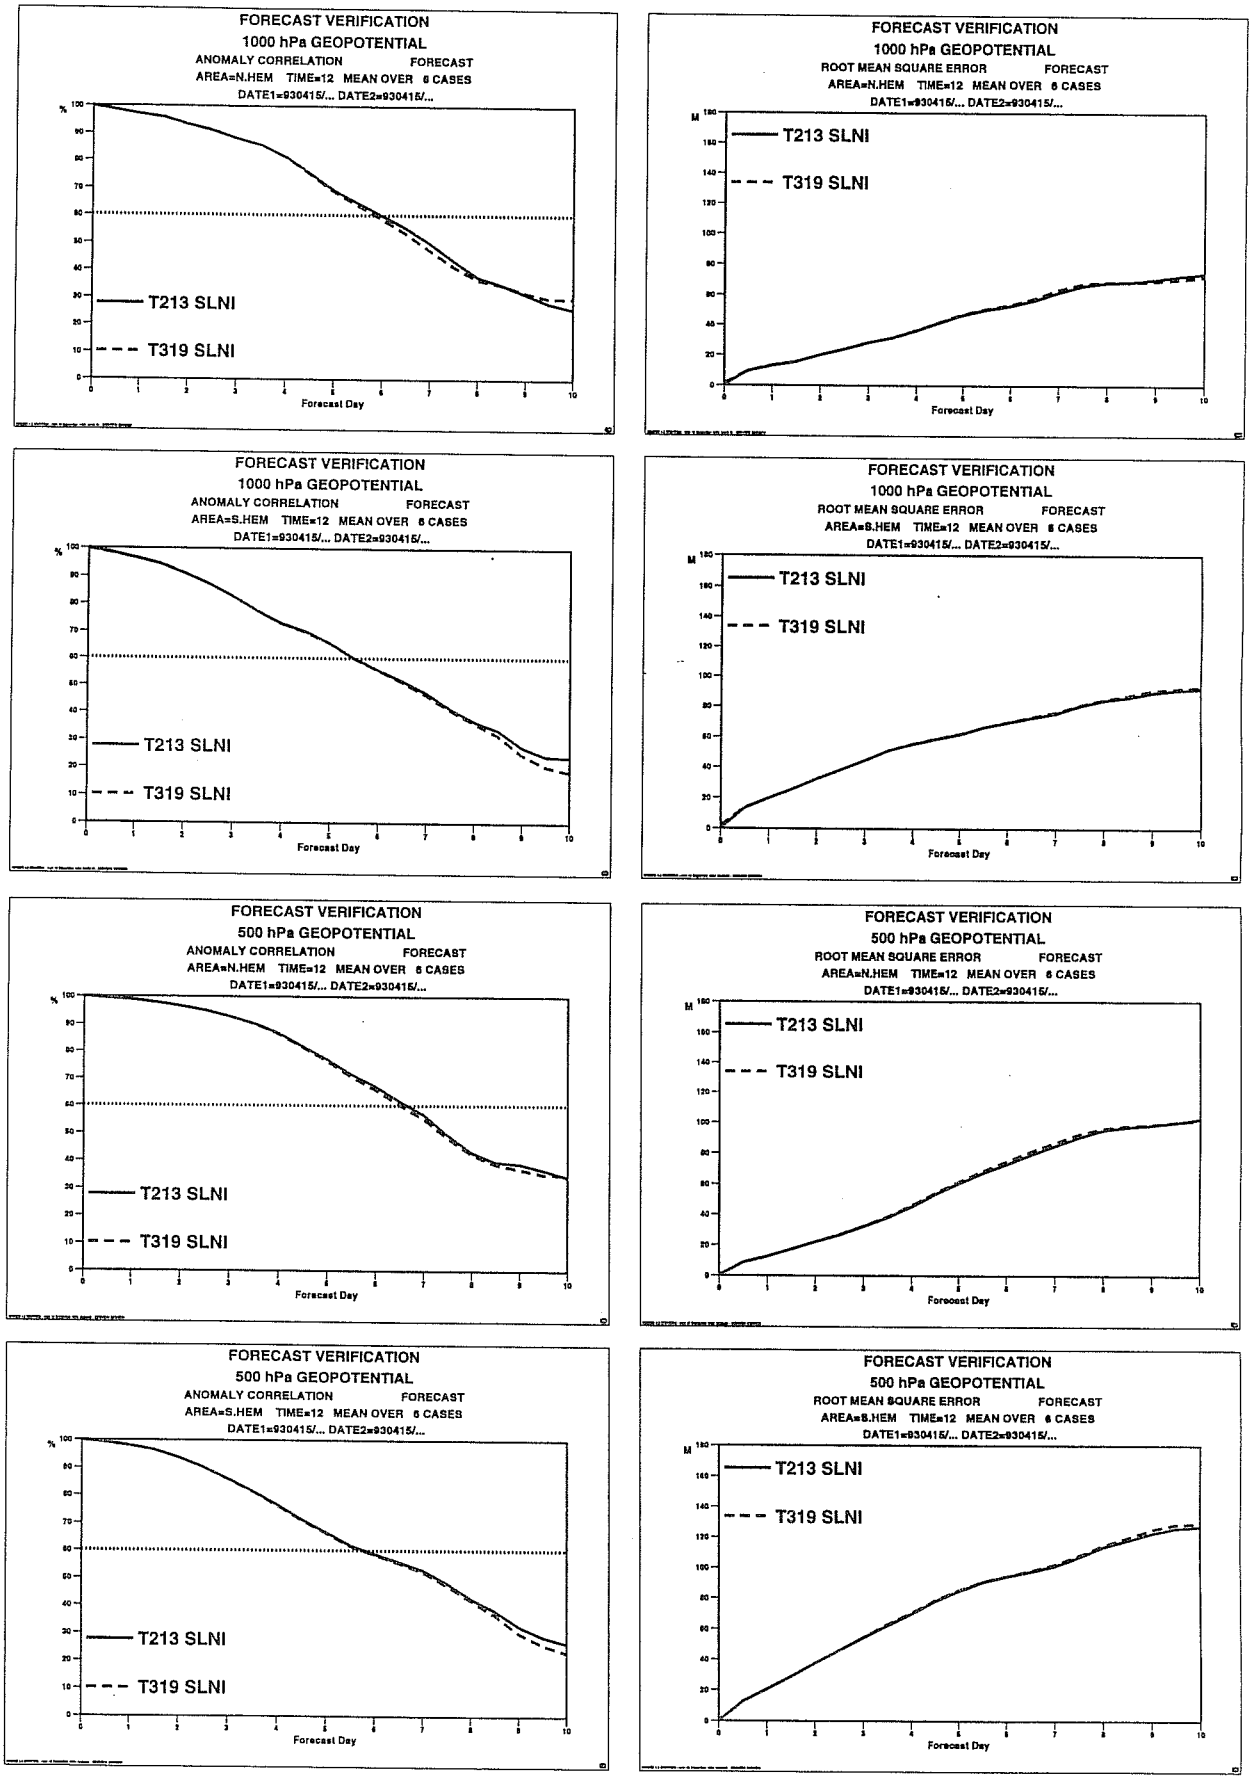

Fig. 13(a): Mean scores of two series of 6 experiments using the same Gaussian grid and the SLNI scheme. The grid at T319 spectral resolution corresponds to the "linear Gaussian grid".

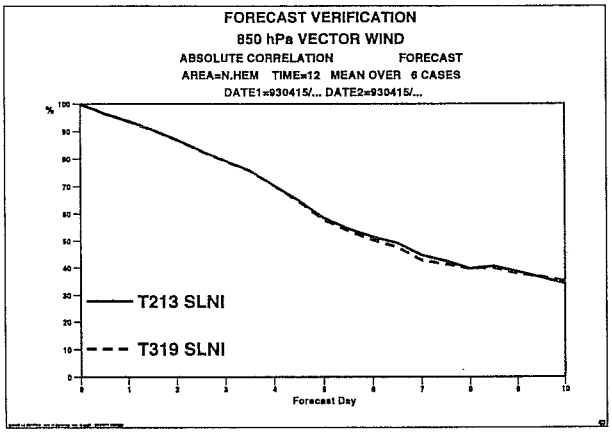

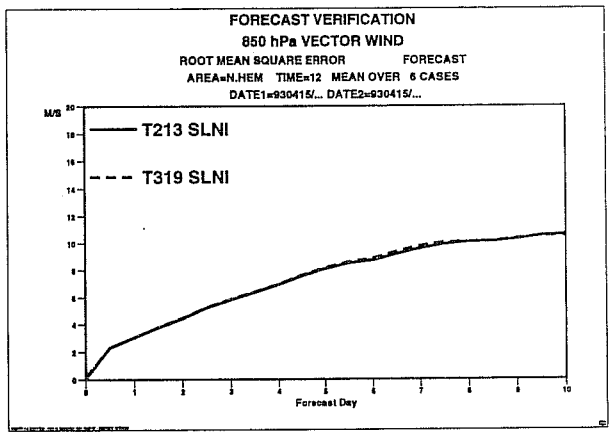

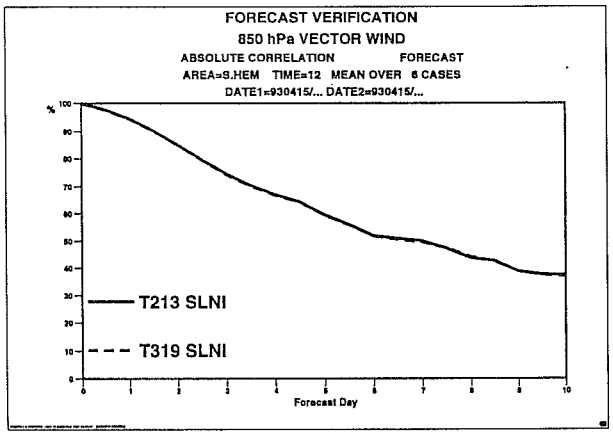

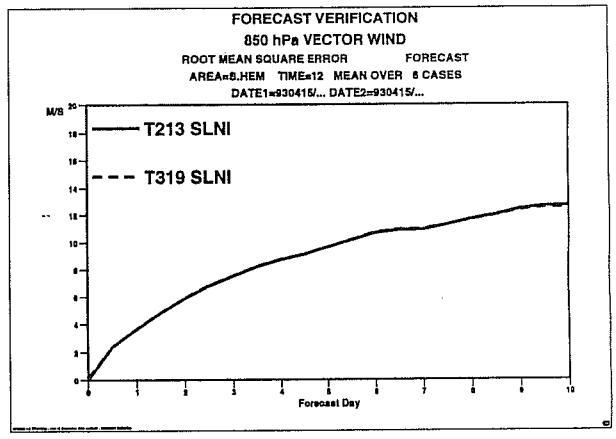

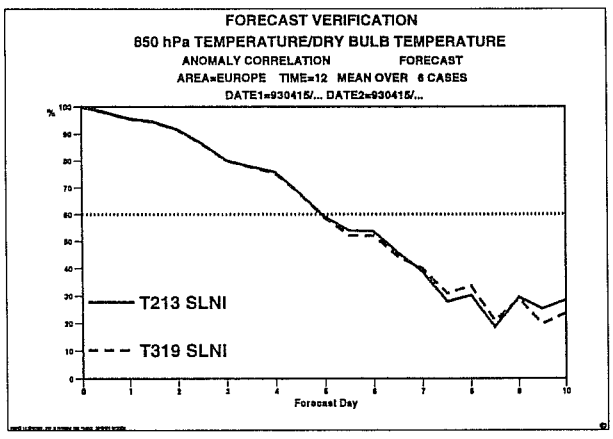

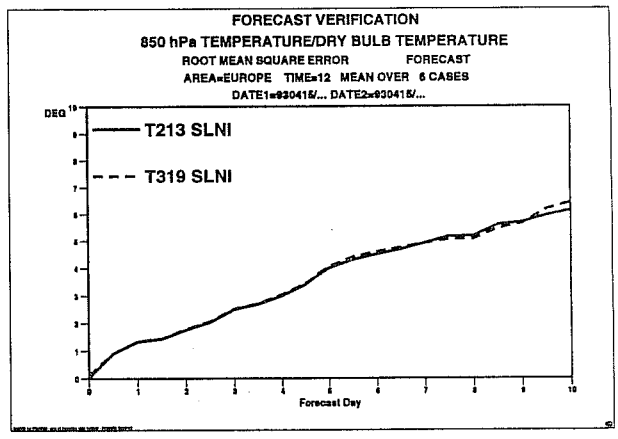

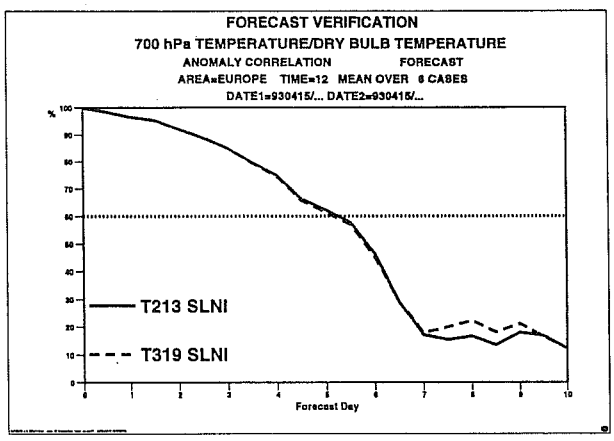

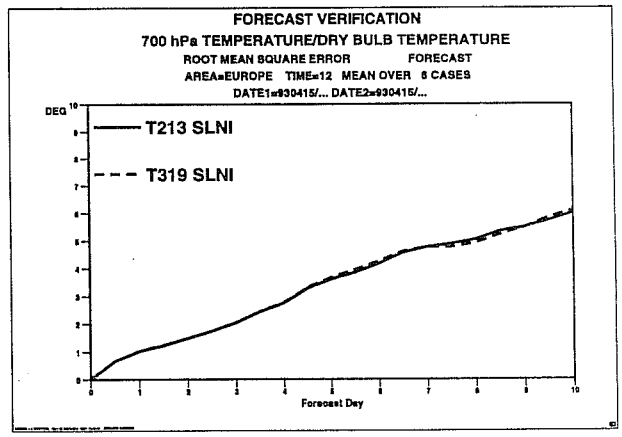

Fig. 13(b): Mean scores of two series of 6 experiments using the same Gaussian grid and the SLNI scheme. The grid at T319 spectral resolution corresponds to the "linear Gaussian grid".

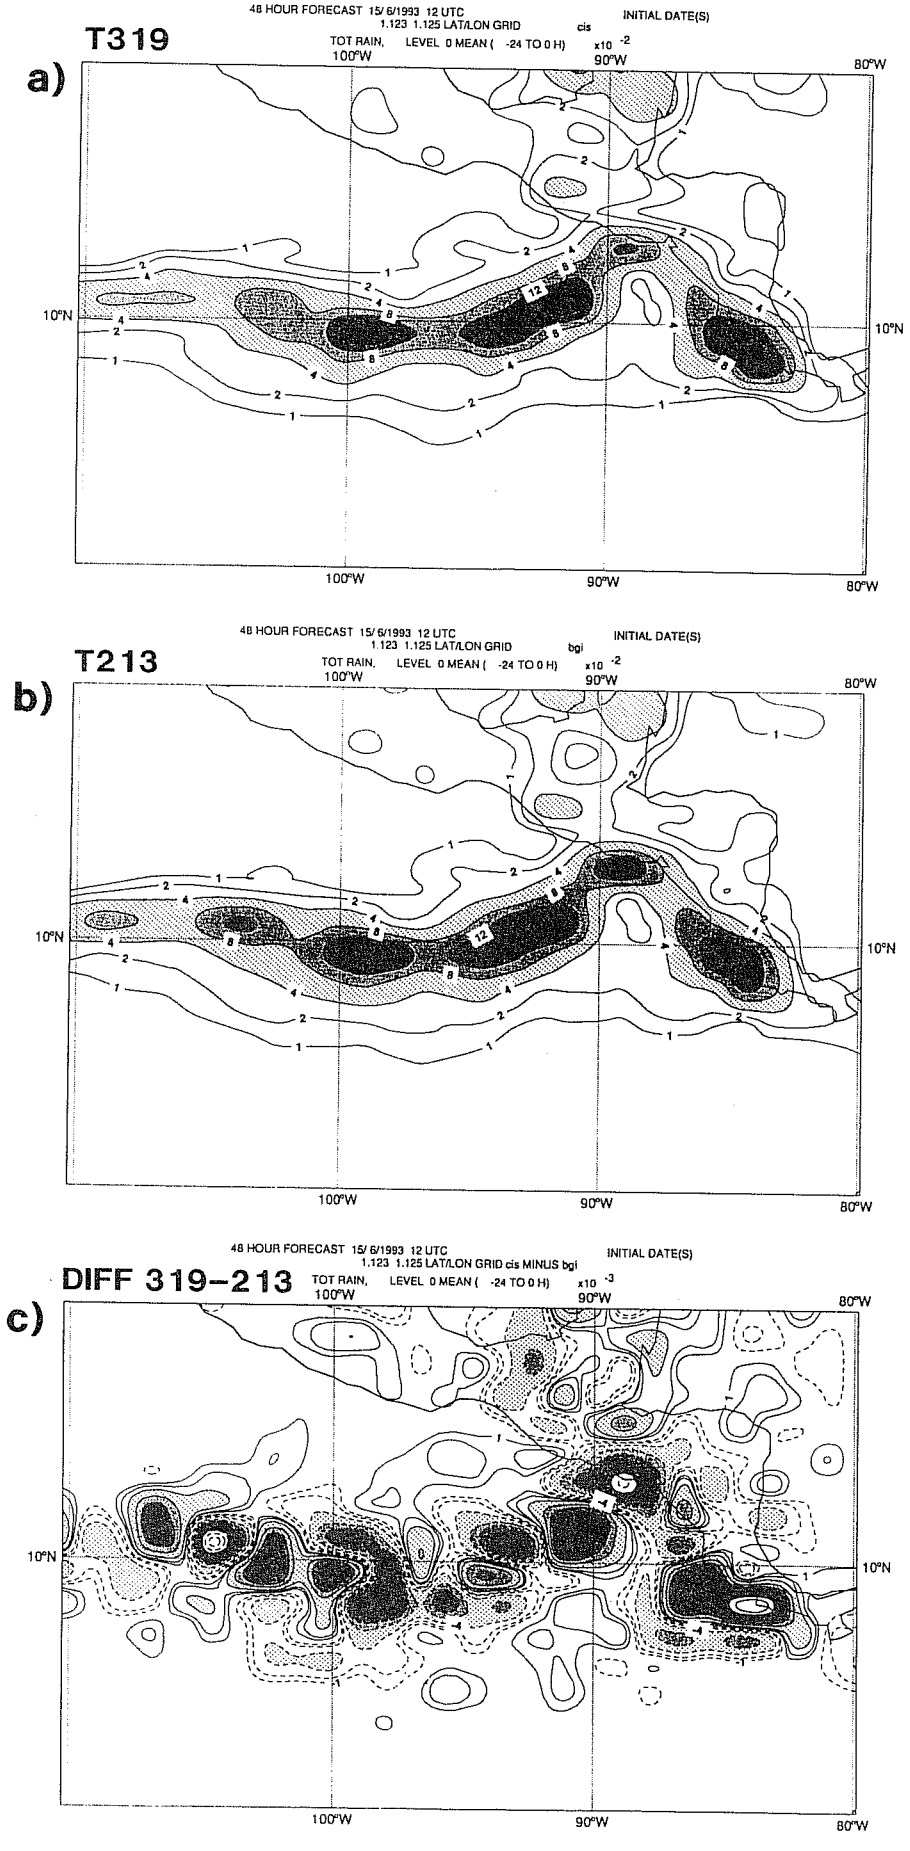

Fig. 14: Maps of precipitation near central America for two experiments using the same Gaussian grid.
(a) experiment at T319 resolution using the linear grid.
(b) experiment at T213 resolution. (c) difference (a)-(b).