Peter Gray

It is with the greatest regret that I have to report that Peter Gray, Head of our Computer Operations Section, died on 7 October 1995. Peter joined the Centre in 1977 as Head of our Operating Systems Section and became the Head of the Operations Section in 1988.

Peter's enthusiasm for life and his enormous personal contribution to both the work and life at the Centre will be sorely missed by all of us. His courageous fight back to health after his recent operation makes his death all the harder to accept and bear.

At Peter's funeral Mike O'Brien, a close friend and colleague of Peter, made a wonderful personal tribute to Peter and I am grateful for Mike's willingness to allow me to share his thoughts on Peter with you.

David Burridge (Director) and staff of ECMWF.

Peter Gray - a personal tribute from a colleague

I worked with Peter for eight years and during that time I learned a great deal from a man who had an abundance of drive and enthusiasm. I had been warned he had quite a bark and that it was often heard, which was very true. He could be forthright about his views - as many of us probably know from having been at the receiving end or overhearing from a distance. His bark was tempered though by an ability to draw people together socially and there were never any hard feelings, when the argument was over it was down to the pub or off to France or just having a cup of coffee. Whenever we had a major shutdown in the computer hall he was always there at three in the morning giving his support and encouragement never asking others to do anything he wouldn't do himself. It sometimes felt as if there was a blonde-haired terrier rushing round at everyone's heels asking how soon we'd be finished. But it was asked in good humour and with the intention of inspiring people to do their best, which was all he ever asked for.

Not only was Pete a superb man to work for he was also a very good friend to many, and to me in particular. At a time in my life when I needed lots of support Pete was there for me and that cemented a friendship I will always cherish.

Kahlil Gibran talks about friendship in his book 'The Prophet' in which he says:

Your friend is your needs answered,
He is your field which you sow with love and reap with thanksgiving
For you come to him with your hunger and you seek him for peace

it then goes on to say

When, you part from your friend grieve not;
for that which you love most in him may be clearer
in his absence, as the mountain to the climber is
clearer from the plain.

I'm finding it very difficult to see across the plain and although it is very painful at the moment I can only be thankful for the friendship that I shared with Pete which will always stay with me.

I feel we have all been honoured to have shared in Pete's very full and active life. He will be remembered for his enormous enthusiasm for living and his ability to draw people together in friendship. I'm sure everyone would join with me in offering sympathy and support to Pete's family at this difficult time and let us all take this opportunity to say farewell to a good friend and spend a quiet moment reflecting upon what was so special about Peter for each of us.

Mike O'Brien

In this issue

This Newsletter is edited and produced by User Support.

The next issue will appear in Autumn 1995.

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, RG2 9AX, UK

Fax: (0118) 986 9450

Telephone: National 0118 949 9000

International ....+44 118 949 9000

Cover

Configuration for Phase 1 of the new data handling system (see page 15).

Editorial

On 4 April 1995, ECMWF introduced a new version of the forecast model that was the culmination of several years research and experimentation. The background to all this is the topic of the article on Page 2.

ECMWF publicises its research workthrough two series of reports, known formally as ECMWF Technical Reports and ECMWF Technical Memoranda. The titles of recently published reports have always been included in the Newsletter, however now we will also begin to publish extended summaries. The first such summary appears on pages 9-12 of this issue.

The fourth meteorological article is a review of the consistency of the ECMWF forecast (page 12).

Of the two articles about the Centre's computer systems, the first is an overview of the new data handling system that will replace, in a phased programme of changes, the existing IBM ES9000 based system. The second article covers a secure access mechanism that is being introduced for Member State users of the ECMWF computer system.

Finally, those with Internet access may like to know that ECMWF now has a world-wide web presence. The URL for our home page is:

http://www.ecmwf.int

Changes to the Operational Forecasting System

Recent changes

No changes which would have had a significant impact on the performance of the ECMWF analysis and forecast system have been introduced during the last three months.

Planned changes

A 3-d variational analysis system will be implemented.

Bernard Strauss

A major operational forecast model change

On 4 April 1995, a substantial and wide-ranging set of IFS modifications was introduced operationally (IFS-CY13R4). This set of model changes was the culmination of several years of development and testing of new parametrization schemes and refinements of the semi-Lagrangian model. This article briefly reviews the substance and rationale for these various changes, and summarizes typical impacts on synoptic/weather parameters. For reference the meteorologically relevant model changes are listed at the end.

It is convenient to discuss those changes involving the numerical formulation separately from the physical parametrization ones, even though there are, of course, fundamental interdependencies. The various components were developed and tested separately and then combined and run in parallel mode before implementation. For conciseness, results from this final combined testing will mostly be used as illustrations.

Numerical formulation

a) Problems in the T213 model

In September 1991, ECMWF began the operational use of a high resolution model with a spectral triangular truncation at 213 wavenumbers and 31 levels in the vertical. This model used a semi-Lagrangian treatment of the advection in which a three-dimensional interpolation was needed to the departure point of each semi-Lagrangian trajectory (scheme SLI). However, a significant increase in the RMS error became apparent, and this was related to an excess of the eddy kinetic energy developing during the 10-day forecast. This could also be seen as an increase in the day-to-day forecast inconsistency. After much study, the semi-Lagrangian treatment was changed in August 1992 to a non-interpolating-in-the-vertical (SLNI) version in which part of the vertical advection is treated in an Eulerian rather than semi-Lagrangian way. This scheme (together with later radiation changes) significantly improved the energetics and levels of consistency.

Nevertheless, the SNLI scheme was not free from problems, the most obvious one being the occurrence of noisy vertical structures when an Eulerian treatment of advection is applied with irregularly spaced model layers.

One hypothesis to explain the improved eddy energy characteristics in the SLNI version compared to the SLI version was that the vertical structures which appear in the former and not in the latter could make the vertical diffusion work much harder, so producing a decrease of the eddy kinetic energy. However, a series of experiments run without vertical diffusion in the free atmosphere, both with the SLNI and with the SLI schemes, indicated that the effect of the vertical diffusion is not very different for the two semi-Lagrangian versions.

Another possibility was that the increase in eddy kinetic energy is produced by over/under-shooting in the dynamical fields either from the spectral truncation of the fields (Gibbs phenomena) or the high order interpolation used in the semi-Lagrangian scheme. The SLNI scheme, which uses only bi-dimensional interpolations, should have less of a problem than the SLI which uses three-dimensional interpolations. Furthermore, the first order only accuracy in space of the Eulerian treatment of the vertical advection in the SLNI scheme could act as a damping mechanism for eddy kinetic energy, therefore compensating somewhat the excess produced by the horizontal interpolations. To test the relevance of the Gibbs phenomena, a series of experiments was run with a Gaussian grid in which the number of degrees of freedom is much closer to the number of degrees of freedom used in spectral space than in the standard Gaussian grid. In the semi-Lagrangian treatment of the advection there is no explicit calculation of quadratic advection terms and the problem of aliassing is not apparent for these terms. Using this 'linear' Gaussian grid, the Gibbs phenomena are greatly reduced but the increase in eddy kinetic energy remained.

If the excess eddy energy was produced by the over/under-shooting in the cubic semi-Lagrangian interpolations, then a shape-preserving or other form of monotonic interpolation should cure the problem. The quasi-monotone scheme proposed by *Bermejo* (1992) (with some improvements) was chosen as it is computationally inexpensive. With this technique, the SLI algorithm produced an evolution of the eddy kinetic energy during the forecast very similar to the one obtained by the SLNI procedure.

Another problem, independent of the advection scheme, was noise in the vertical velocity field at the lowest levels of the model in the vicinity of even quite small orographic features. Present in almost all forecasts, independently of the synoptic situation, it was particularly marked over the sea near Taiwan or in a circular pattern around Hawaii. An attempt was made to reduce this noise by increasing the horizontal diffusion in the model but this actually increased it.

Various tests were carried out which demonstrated that the problem of low-level noise in vertical velocity near orographic features was not due to some form of aliassing.

However, inspection of the spectra of vorticity and divergence at the lowest model level of a T213 forecast after 24 hours of integration shows that the shape of the high wave-number part of both fields is clearly different, Fig 1. The one for divergence is 'raised' relative to that of vorticity. 'Lowering' the high wavenumber tail of the spectral representation of the orography through the application of a high (8th) order diffusion operator resulted in a close matching of the tails of the spectra of vorticity and divergence. The corresponding low-level vertical velocity fields exhibited none of the noise seen previously.

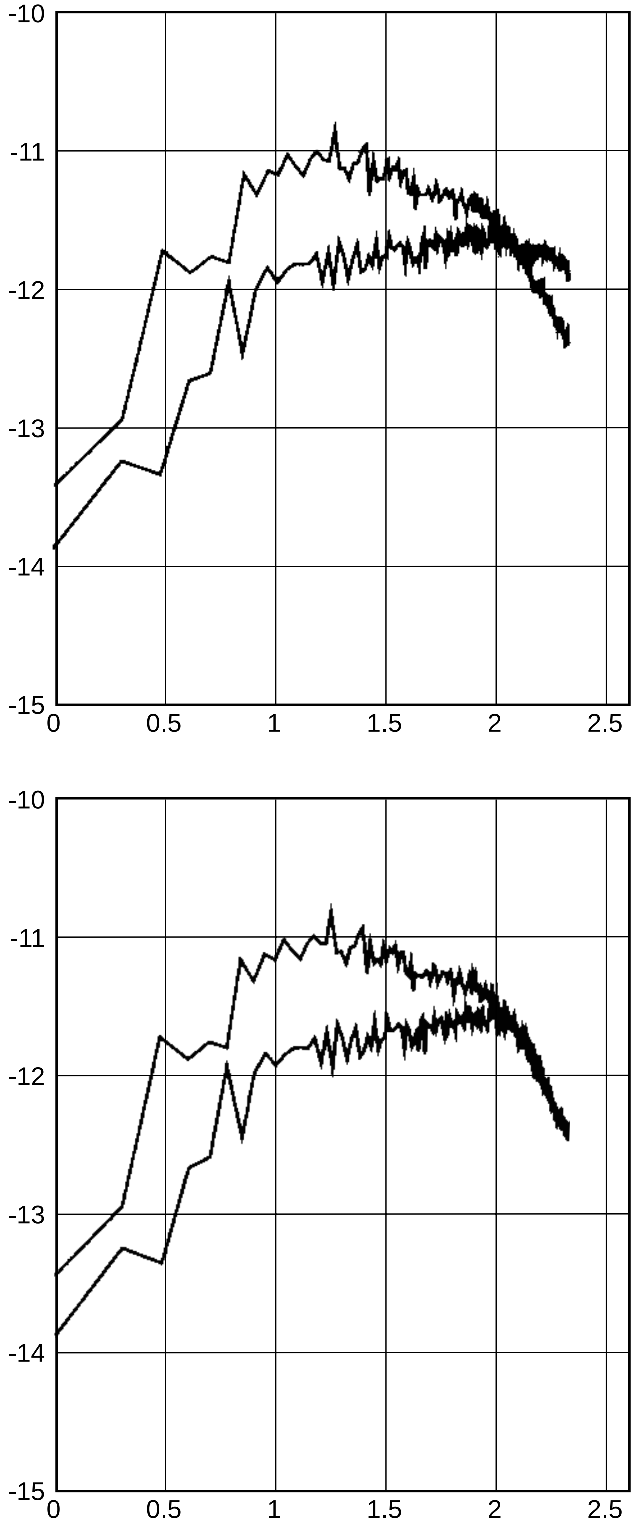

Fig. 1 Spectra of vorticity and divergence at model level 31 of a T213 forecast. Top is unsmoothed and bottom with smoothed mean orography.

More details of the above experimentation can be found in *Hortal* (1994).

b) Changes

As a result of the above diagnostics, it was proposed switching back to the SLI version of the model in operations but using quasi-monotone interpolations. However, when applied to the vertical interpolation of the temperature field, this produced a 'smoothing' of the tropopause leading to a large warm bias. Similarly, the application of the quasi-monotone procedure in the vertical interpolation of the wind fields led to large errors at the position of the stratospheric maximum of the tropical zonal wind (QBO). Since the largest contribution to the improved eddy kinetic energy statistics came from the quasi-monotone procedure in the horizontal, a fully interpolating semi-Lagrangian scheme, using the quasi-monotone procedure in the horizontal part of the interpolations only was chosen.

A further change is that the humidity field is no longer transformed to spectral space, therefore avoiding the introduction of Gibbs phenomena. The advection is handled by the semi-Lagrangian procedure with quasi-monotone cubic interpolation in three dimensions. This ensures that no unphysical negative value of the humidity arises, and the need to correct for negative humidities can only occur from the integration of the moist parametrizations.

Finally, an advective treatment of the Coriolis term in the momentum equation, first proposed by Rochas at Météo-France, has been introduced. In this treatment, the term is $2\vec{\Omega}\wedge\vec{v}_h$ is expressed as $\frac{d}{dt}\left(2\vec{\Omega}\wedge\vec{r}\right)$ and added to the total time derivative of the horizontal wind.

The new prognostic cloud scheme

An important change in cycle 13R4 is the introduction of the new prognostic cloud scheme. The scheme was developed by *Tiedtke* (1993) and is based on two prognostic equations for cloud liquid water/ice and cloud fraction, respectively. Subsequent modifications and testing are summarised in *Jakob* (1994). The scheme is process-oriented, i.e. source and sink terms are linked to processes known to produce or dissipate clouds. The comprehensive nature of these component processes is summarised in Fig. 2 and discussed in the following paragraphs.

The generation processes include condensation due to large-scale ascent, cumulus convection, boundary layer turbulence and radiative cooling. Cloud dissipation is accounted for through evaporation due to large-scale descent, cumulus-induced subsidence, radiative heating and turbulence at both cloud tops and sides, as well as through precipitation processes. The schemes thus give strong linkages between the physical and dynamical processes in the model.

The link to cumulus convection is achieved by considering convective clouds as being produced by the detrainment of updraught condensates into the environmental air. This description of convective clouds is an important part of the cloud scheme, and, by representing stratiform clouds produced by convection, forms a substantial extension of the cumulus convection parametrization as well. It leads to a much more realistic description of the different types of convective clouds such as precipitating and non-precipitating cumuli, cumulonimbus clouds, anvils and cirrus debris from convective processes.

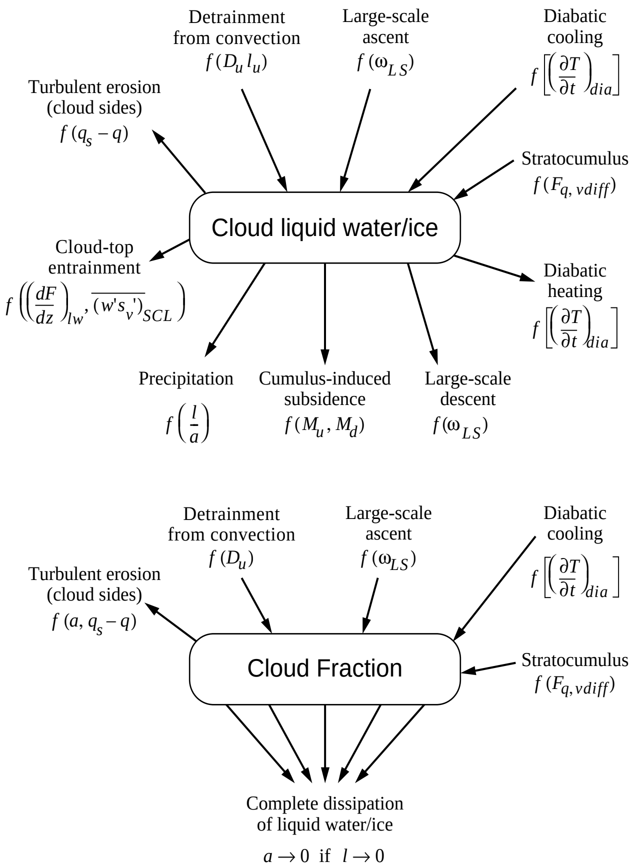

Fig. 2 A summary of the processes incorporated in the prognostic cloud scheme.

Stratocumulus clouds are also specifically included in the scheme. They are represented through the linkage of the scheme to the boundary layer moisture flux produced by the vertical diffusion scheme. Cloud top entrainment processes are taken into account as a function of buoyancy in the mixed layer.

The formation of stratiform clouds is determined by the rate at which the saturation specific humidity decreases due to upward vertical motion and radiative cooling.

Evaporation processes are accounted for in several ways. A main contribution to evaporation is determined by the rate of increase of saturation specific humidity caused by large-scale and cumulus-induced subsidence and radiative heating. Turbulent processes lead to evaporation at the cloud sides proportional to the saturation deficit in the grid-box. Turbulent evaporation at cloud tops (entrainment) is represented through a dependency on the radiative cooling rate.

Precipitation processes use simplified microphysical formulations; these relate the generation of precipitation to the local cloud water/ice content and also take into account the main precipitation enhancement processes such as the Bergeron-Findeisen mechanism, collision and coalescence.

The scheme described above is unique in treating the main cloud-related processes in a consistent way by forecasting both cloud fraction and cloud water/ice content with prognostic equations. The strong coupling of the cloud to the other parametrized processes provides a potentially powerful diagnostic capability for further refinement of the overall model physics.

The representation of orographic effects

The ECMWF model has used an 'envelope' orography since 1983 whereby the grid square mean height is enhanced by an amount proportional to the standard deviation of the sub-grid scale heights. This has provided positive impact on the forecast dynamics at the expense of significant loss of low-level data availability due to the mismatch between model and actual station altitudes. Additional problems associated with the use of an envelope orography have been identified; these include a contribution to warm continental biases in forecast temperatures, overprediction of convective precipitation and the spurious broadening and intensification of heavy rain events associated with orographic forcing. Consequently these problems have been met by reverting to a smoothed mean orography (see earlier) together with a new sub-grid scale orography parametrization scheme to combine, and improve on, the benefits of both orographies.

a) Sub-grid scale orography scheme

The new scheme is designed to represent nonlinear low-level mountain drag together with the already familiar 'gravity wave drag' due to the propagation and dissipation of orographically-excited waves. An important and novel part of the scheme is that it explicitly represents the blocking of low-level flow and the associated form drag due to flow separation caused by sub-grid scale orography that is assumed to intersect the model levels. The depth over which this drag is parametrized is determined by the height of the sub-grid orography and an inverse Froude number characterising the flow at any given location and time. Flow above this 'blocked' layer goes over the orography and can generate gravity waves. Extensive use of PYREX data has been made to examine these ideas, to adjust the parameters of the scheme and to verify T106 and T213 forecasts of drag and momentum profiles above the Pyrenees.

Orography data is provided in the form of four grid-point fields describing the height, orientation, anisotropy and slope of the sub-grid scale orography together with the grid scale mean heights. The combination of ***mean*** orography and this new scheme has been shown to be equal or superior to that of envelope orography plus the old gravity wave drag scheme while no longer suffering any disadvantages of envelope orography.

Since the new scheme influences the flow well above the mean orographic height, it is possible to derive an 'effective' orography increment generated by the scheme which can be interpreted as additional to the mean resolved orography. An example, computed as a ten-day forecast mean, is shown in Fig. 3 together with the basic mean orography. Complete details of the scheme, its development and testing can be found in *Lott and Miller* (1995) and *Baines and Palmer* (1990).

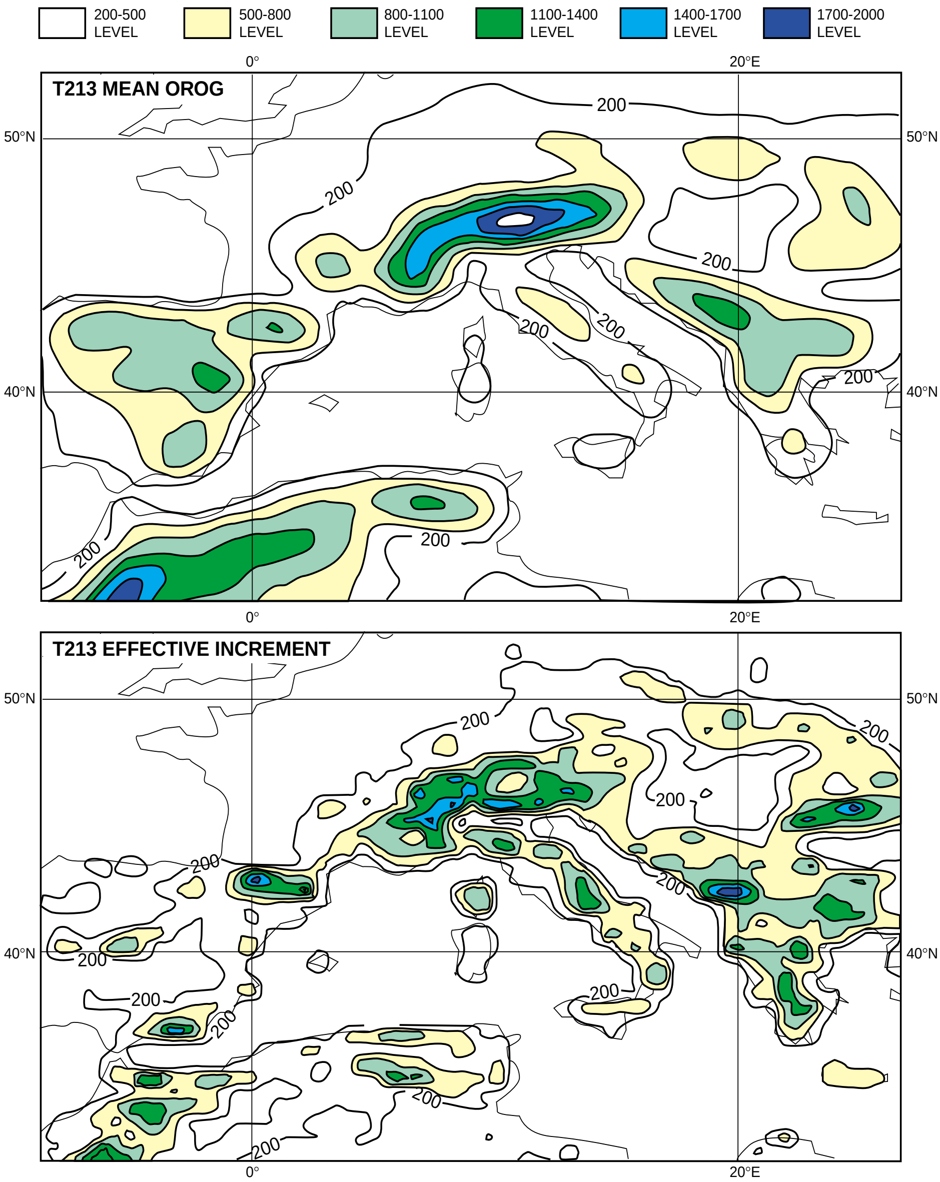

Fig. 3 The smoothed mean orography for the T213 model (top), and (bottom) the 'effective' T213 orography increment implied by the new subgridscale orography scheme. (Computed as a ten-day forecast mean, initial date 15/1/95.)

Summary of results

Prior to operational testing, two periods of approximately three weeks each (initial dates 1/6/94 and 6/12/94) were run with the full IFS. A selection of results follows which is chosen to illustrate the benefits of the new cycle in a variety of ways. Fig. 4 shows objective scores for the two E-suites referred to above, and it is encouraging to note improvements in skill for all areas. For the June period the comparison is made against the modified operations that included the soil moisture analysis described in the ECMWF newsletter (No. 69).

The new cycle has major impacts directly on the forecast cloud distribution and amount, and indirectly on near-surface parameters such as two-metre temperature. Although these are predominantly from the cloud changes, the other changes also contribute quite significantly.

Fig. 5 compares the total cloud cover for the mean of the twenty five-day forecasts made with the model operational in June 1994 (with the diagnostic cloud scheme) with the total cloud cover for the new cycle. Immediately apparent is the increased cloudiness particularly in the midlatitude storm tracks. There is also more cloud over the subtropical oceans (excessively so) and less over the subtropical desert areas. Although there is some improvement in the marine stratocumulus areas, these do not extend close to the coasts.

The impact of the new cycle on the cloud cover and two-metre temperature forecasts for Europe is shown in Fig. 6 which compares several forecast ranges against SYNOP observations.

Previous negative cloud biases of more than one okta before April 1995 have been reduced by about 50%, while the large daytime warm biases discussed in the previous newsletter have been substantially alleviated.

As an example of the impact of the reduced numerical 'noise' discussed earlier, Fig. 7 compares the 48-72 hr forecast precipitation fields (reflecting the smoother vertical velocities).

Operational experience since April 1995 has supported the above examples and the impacts both on skill scores and weather parameters have been generally good. This experience has also confirmed the improvements in precipitation distribution and amount particularly in the vicinity of mountain regions.

Such a major model change will inevitably exhibit new problems as well as solving or alleviating pre-existing ones. However, the current model should provide an effective foundation for further development and refinement in the next few years.

Martin Miller, Mariano Hortal, Christian Jakob

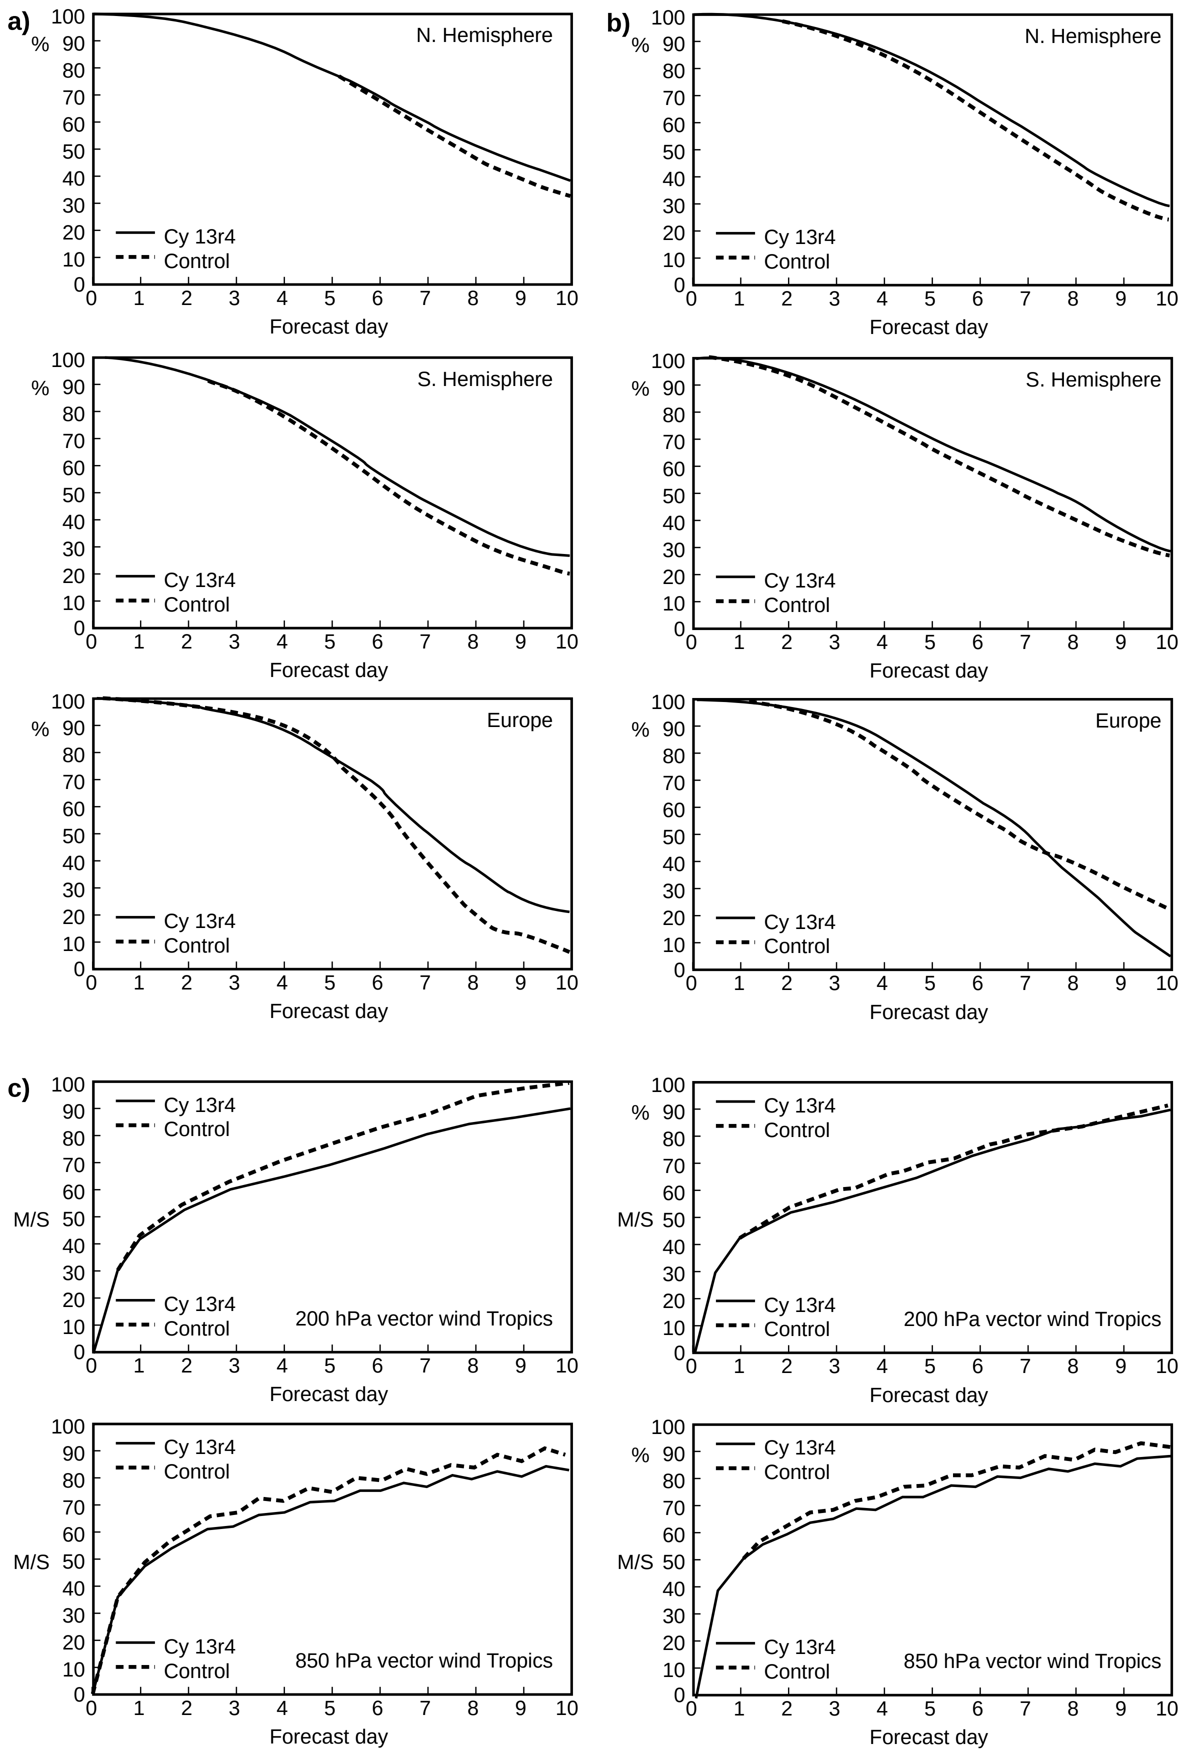

Fig. 4 a) 500 mb height anomaly correlations for the December E-suite (21 forecasts). b) As a) for the June E-suite (20 forecasts). c) Tropical wind RMS scores. LHS from June, RHS from December.

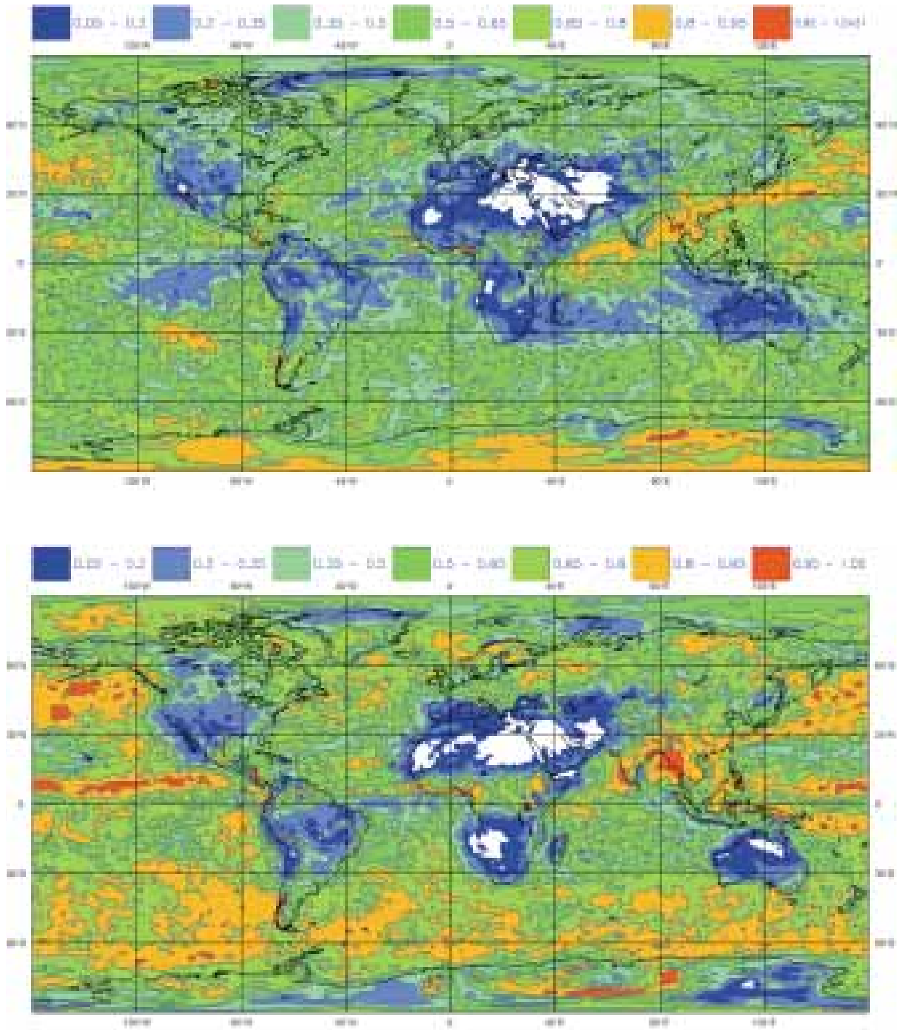

Fig. 5 Upper: Total cloud cover averaged over all (D+5) operational forecasts for the June period (1/6/94-20/6/94). Lower: As above but for the new cycle.

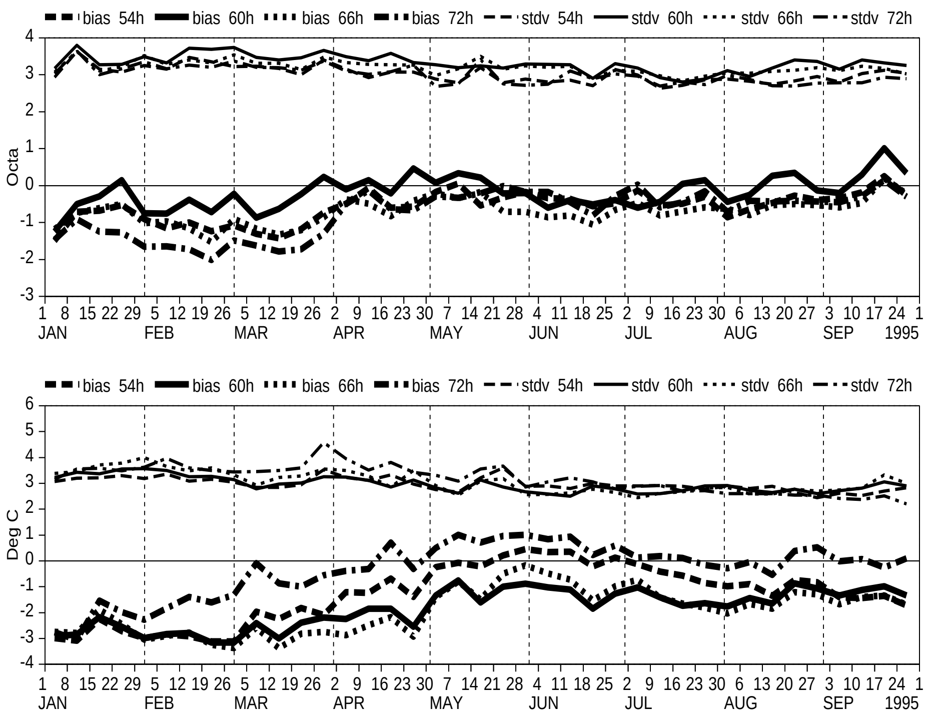

Fig. 6 A weekly time series of the biases of forecast cloud cover and two-metre temperature as calculated against European SYNOP reports.

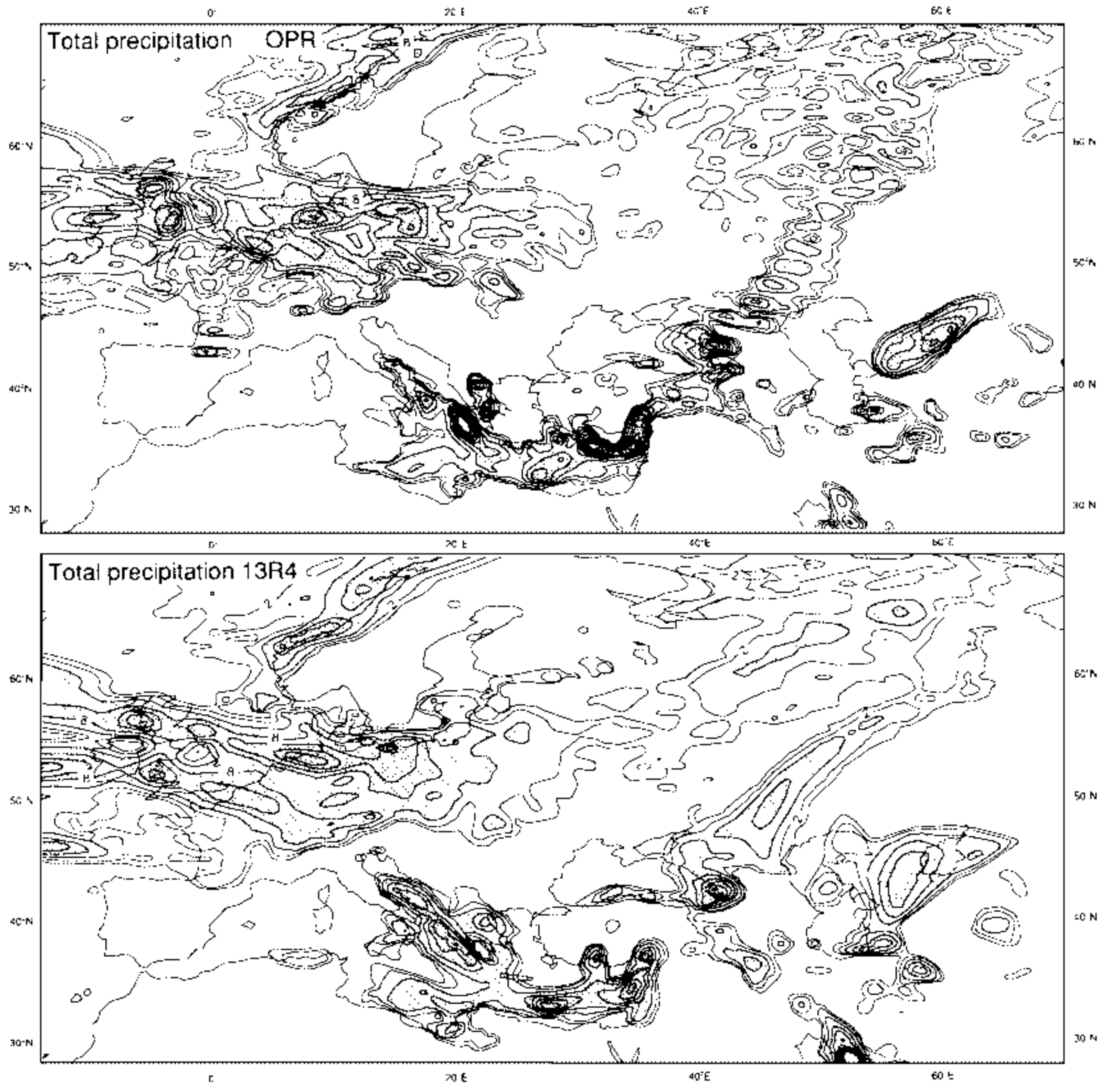

Fig. 7 Forecast model 24 hr precipitation fields for the 48-72 hr forecast. Upper: Operations - initial date 5/11/94. Lower: New cycle.

Annex

Meteorologically relevant model changes:

- prognostic cloud scheme
- smoothed mean orography
- new sub-grid scale orography parametrization
- zenith-angle dependent ocean surface albedo
- fully interpolating semi-Lagrangian scheme
- averaging of the RHS of all equations along the trajectories using linear interpolation at the departure points for these terms
- advective form of the Coriolis terms
- quasi-monotone cubic interpolation in the horizontal for the (t-*D*t) terms of all the equations and likewise in the vertical for moisture (q)
- use of (RT) rather than T as spectral variable
- grid-point specific humidity
- modified reduced Gaussian grid

References

[1] Baines, P G and T N Palmer, 1990: Rationale for a new physically-based parametrization of sub-grid scale orographic effects. *ECMWF Tech Memo 169.*

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

[3] Hortal, M, 1994: Recent studies of semi-Lagrangian advection at *ECMWF. ECMWF Tech Memo 204.*

[4] Jakob, C, 1994: The impact of the new cloud scheme on ECMWF's integrated forecasting system (IFS), Proceedings of ECMWF/GEWEX workshop on modelling, validation and assimilation of clouds, *ECMWF,* November 1994.

[5] Lott, F and M Miller, 1995: A new sub-grid scale orographic drag parametrization: its formulation and testing. Submitted to *QJRMS.*

[6] Tiedtke, M, 1993: Representation of clouds in large-scale models. *Mon Wea Rev,* **121,** 3040-3061.

Summary of ECMWF Technical Memorandum No. 214

The anomalous rainfall over the USA during July 1993: sensitivity to land surface parametrization and soil moisture anomalies

A C M Beljaars, P Viterbo, M J Miller, A K Betts

July 1993 showed anomalously high precipitation over the central USA with exceptional flooding of the Mississippi (*Kunkel et al*, 1994). During this month, a new model version (cycle 48, hereafter called CY48) of the ECMWF model with the land surface scheme described by *Viterbo and Beljaars* (1995), ran in parallel with the operational forecast system (model cycle 47; hereafter called CY47). Such parallel runs are used for final testing of system changes prior to operational implementation. The parallel run included data assimilation and 10 day forecasts with the new model version at operational resolution (T213L31). Figure 1 shows a comparison of the CY47 model and CY48 24 hour precipitation forecasts for day three. It is evident that CY48 produces more realistic precipitation fields.

Although the two model versions have a number of differences, it is believed that the revised land surface scheme is the main contributor to the improved precipitation forecasts. The new scheme does not specify a climatological deep boundary condition for soil moisture and therefore has a much better potential for handling soil moisture anomalies. To illustrate this aspect, sensitivity experiments were also carried out with the new CY48 at T106L31 resolution. An ensemble of three 30 day integrations were initialized with soil moisture at field capacity and compared with a similar ensemble initialized at 25% of soil moisture availability. The 'wet' integrations produce much more precipitation than the 'dry' integrations, showing that evaporation influences precipitation considerably. These results suggest that there is persistence in the behaviour of precipitation and so there may be seasonal predictability potential related to the depletion time scale of the soil moisture reservoir (*Rowntree and Bolton*, 1983; *Mintz*, 1984).

The treatment of land surface processes in large scale atmospheric models has always been a difficult problem, because of the complicated interactions between atmosphere, vegetation and soil, and the lack of data for different climatological regimes. Current physically based land surface schemes may model atmosphere surface interaction with Monin Obukhov similarity, treat bare soil evaporation as controlled by soil moisture in a shallow top soil layer, treat evaporation from vegetation with a stomatal resistance controlled by soil moisture in the root zone, they may represent potential evaporation from snow and/or an interception reservoir, and they may also represent some sort of soil hydrology (e.g. *Dickinson et al*, 1986; *Sellers et al*, 1986; *Noilhan and Planton*, 1989; *Abramopoulos et al*, 1988; *Blondin*, 1991; *Viterbo and Beljaars*, 1995). These models can vary widely in their implementation depending on the details of the model and the level of complication that is adopted (e.g. vertical resolution in the soil, number of geographically dependent terrain parameters etc.). Although most schemes have been validated and calibrated with the help of field data, the differences between individual models can still be large, as is becoming increasingly clear from the Project for Intercomparison of Land-surface Parametrization Schemes (PILPS; *Henderson-Sellers et al*, 1993). The behaviour of land surface schemes on the seasonal time scales is less well documented and less well tested than their behaviour on the diurnal time scale.

The impact of land surface processes in General Circulation Models (GCM's) has been studied by *Rowntree and Bolton* (1983), *Dickinson and Henderson-Sellers* (1988), *Nobre et al* (1991), *Lean and Rowntree* (1993), *Atlas et al* (1993) and *Milly and Dunne* (1994) (and reviewed recently by *Garratt*, 1993). *Mintz* (1984) concluded that a coupling exists between evaporation and precipitation in the sense that precipitation increases over land when evaporation increases. This implies a positive feedback from the recirculation of precipitation through the soil moisture reservoir, which may lead to prolonged persistence of anomalous wet or dry spells. Such a persistence was indeed found by *Oglesby* (1991) from model simulations, and suggested in a data study by *Namias* (1958), in which significant lagged correlations were found in the 700 hPa field over the Midwest of the USA between spring and summer, and from a correlation of spring precipitation anomalies and summer temperature anomalies.

Although the precipitation response to evaporation is marked, it is not very clear how this response works in detail. *Rowntree and Bolton* (1983) present the general notion that evaporation increases the moisture content of the troposphere and brings the air closer to saturation and therefore facilitates precipitation. However, given a certain incoming net radiation, soil moisture mainly affects the partitioning of sensible and latent heat flux at the surface. Consequently the equivalent potential temperature (θ_e, determining the condition for convective precipitation) in the boundary layer is not affected by the Bowen ratio at the surface. *Betts et al* (1994) and *Betts and Ball* (1995) demonstrate with observational data and a simple boundary layer model that reduced evaporation, and therefore increased heating at the surface, increases entrainment at the top of the boundary layer and thus increases entrainment of low θ_e air from above the boundary layer. This leads to lower θ_e in the boundary layer, a more stable troposphere and might therefore lead to less convective precipitation.

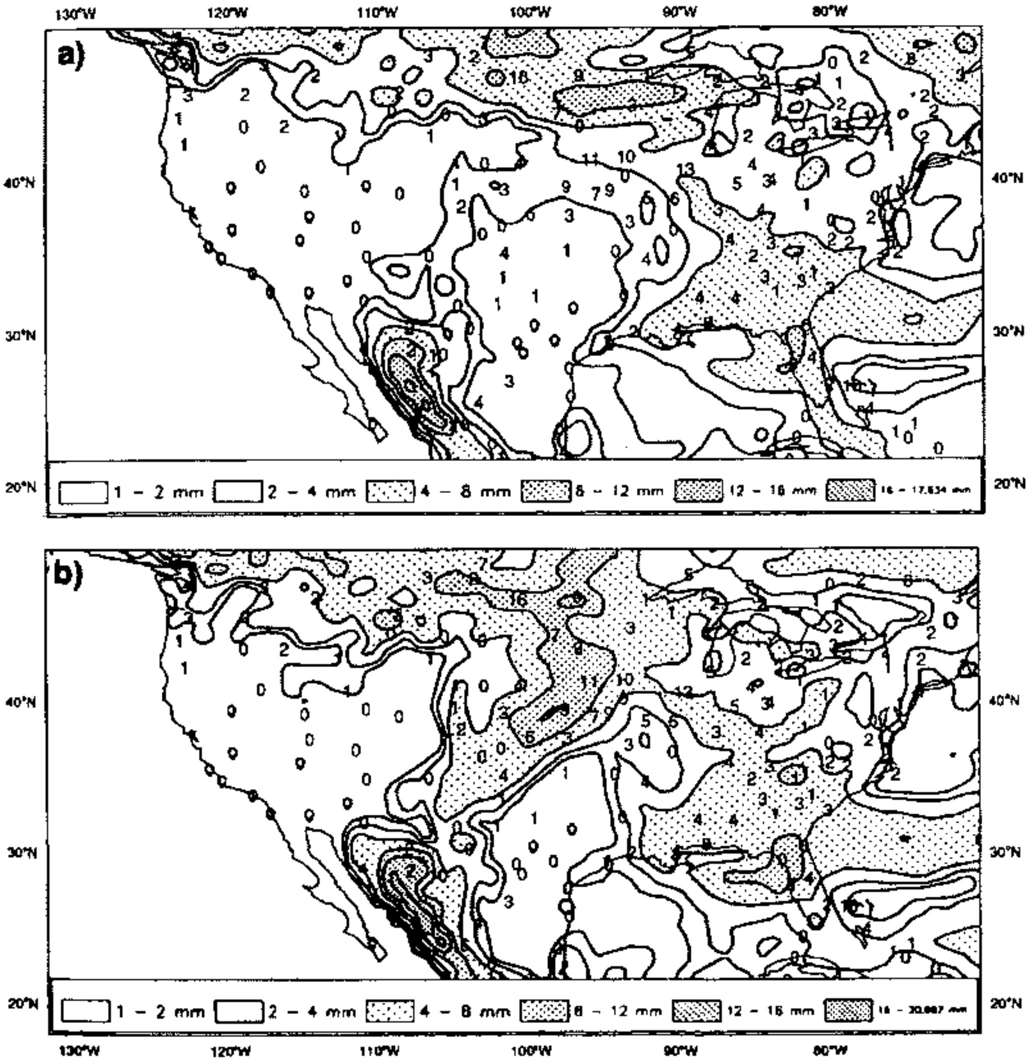

Fig. 1 Mean forecast precipitation of all 48 to 72 hour forecasts verifying between 9 and 25 July, with CY47 (A) and with CY48 (B). The contours are at 1, 2, 4, 8 .... mm/day. The printed numbers are station observations in mm/day.

Summer precipitation over the Mississippi basin is closely linked to a persistent flow pattern, sometimes referred to as the USA summer monsoon. Moisture is transported from the Gulf of Mexico in a northward boundary layer flow over Mexico, Texas, and Oklahoma, curving gradually eastward over the plains (*Rasmusson*, 1967, 1968, 1971). Severe storms triggered by upper air disturbances coming from the West, create much of the summer precipitation.

Benjamin and Carlson (1986) and *Lanicci et al* (1987) discuss the role of differential advection, with a warm/dry southwesterly flow from the Mexican plateau overlying a moist southerly flow from the Gulf of Mexico, giving a capping inversion which allows the build up of large conditional instability. As a consequence, when the soil is more moist over the Mexican Plateau the heating is smaller, the stabilizing inversion is weaker, and convective precipitation is not inhibited.

In *Beljaars et al* (1995) results were presented for July 1993 for the two forecast systems that were running in parallel at ECMWF. The operational system (CY47) has a land surface scheme that is heavily constrained by the deep soil climatological boundary condition for soil moisture and temperature. The new system (CY48) produces its own soil moisture by integrating the soil variables with input from the atmospheric model (e.g. precipitation and evaporation) in the short range. In this way it has the potential to simulate soil moisture anomalies. The new system also has a different formulation for the mixed layer in which a dry entrainment formulation is used. The pre-operational parallel run of CY47 and CY48 showed that the anomalous precipitation over the centre of the USA was captured much more realistically in the 72 hour time range with the new model version. The area with large precipitation differences has only small differences in evaporation, indicating that the local moisture supply is not the controlling factor. Inspection of the thermodynamic profiles indicates unrealistically strong inversions with CY47, capping the boundary layer moisture, and inhibiting the development of convective precipitation. It is deduced that these unrealistic inversion structures develop as the result of excessive heating upstream of the area of interest. The new model version produces less heating because more moisture is available for evaporation. In this way evaporation affects precipitation in the downstream area, similar to the mechanism proposed by *Benjamin and Carlson* (1986) and *Lanicci et al* (1987). Also *Rowntree and Bolton* (1983) observe downstream propagation of moisture anomalies in their simulations for the European area.

The second part by Beljaars et al (1995) shows the result of two ensembles of 30 day integrations for July 1993 using the model CY48. The purpose of these experiments is to study the ability of the model to maintain soil moisture. Soil moisture in the first ensemble is initialized at field capacity, whereas the second ensemble is initialized at 25% availability. The wet integrations show much more realistic precipitation for the month of July, but now the evaporation is also quite different. On these longer time scales precipitation and evaporation are more closely in balance locally, and it is more difficult to separate cause and effect. However, the mechanism as diagnosed in the first part of the paper, namely that the difference in precipitation in the central USA is linked to a difference in thermodynamic structure associated with evaporation differences upstream is equally valid in these longer timescales.

Irrespective of the detailed mechanism responsible (which is hard to determine in 30-day ensemble means), these sensitivity experiments indicate a pronounced feedback between land surface hydrology and precipitation on large space and long time scales, associated with the memory of the land surface boundary condition. There is clear indication of predictive skill in the monthly timescale resulting from the long time constant of the soil moisture reservoir. This result underlines the importance of the observation and analysis of soil moisture (e.g. *Bouttier et al*, 1993).

References

1 Abramopoulos, F, C Rosenzweig and B Choudhury, 1988: Improved ground hydrology calculations for global models (GCM's): soil water movement and evapotranspiration. *J Climate,* **1,** 921-941.

2 Atlas, R, N Wolfson and J Terry, 1993: The effect of SST and soil moisture anomalies on GLA model simulations of the 1988 US summer drought. *J Climate,* **6**, 2034-2048.

3 Beljaars, A C M, P Viterbo, M J Miller and A K Betts, 1995: The anomolous rainfall over the USA during July 1993: Sensitivity to land surface parametrization and soil moisture anomalies. *ECMWF Technical Memorandum No. 214.*

4 Benjamin, S G and T N Carlson, 1986: Some effects of surface heating and topography on the regional severe storm environment. Part I: Three-dimensional simulations. *Mon Wea Rev,* **114,** 307-329.

5 Betts, A K and J H Ball, 1995: The FIFE surface diurnal cycle climate. *J Geophys Res,* in press.

6 Betts, A K, J H Ball, A C M Beljaars, M J Miller and P Viterbo, 1994: Coupling between land-surface boundary-layer parametrizations and rainfall on local and regional scales: Lessons from the wet summer of 1993. Fifth Conference on global change studies, Amer Meteor Soc, 74th Ann Meeting, Nashville, Tenn, 174-181.

7 Blondin, C 1991: Parametrization of land-surface processes in numerical weather prediction. Land surface evaporation - Measurement and parametrization, T J Schmugge and J-C Andre, Eds, Springer-Verlag, 31-54.

8 Bouttier, F, J-F Mahfouf and J Noilhan, 1993: Sequential assimilation of soil moisture from atmospheric low-level parameters, Part II: Implementation in a mesoscale model. *J Appl Meteor,* **32,** 1352-1364.

9 Dickinson, R E, A Henderson-Sellers, P J Kennedy and M F Wilson, 1986: Biosphere atmosphere transfer scheme (BATS) for the NCAR community climate model. *NCAR Tech Note, NCAR/TN-275+STR,* 69 pp.

10 Dickinson, R E and A Henderson-Sellers, 1988: Modelling tropical deforestation: A study of GCM land-surface parametrizations. *Quart J Roy Meteor Soc,* **114,** 439-462.

11 Garratt, J R, 1993: Sensitivity of climate simulations to land-surface an atmospheric boundary-layer treatments - A review. *J Climate,* **6,** 419-449.

12 Henderson-Sellers, A, Z-L Yang and R E Dickinson, 1993: The project for intercomparison of land-surface parametrization schemes. *Bull Amer Meteor Soc,* **74,** 1335-1349.

13 Kunkel, K E, S A Changnon and J R Angel, 1994: Climate aspects of the Mississippi river basin flood. *Bull Amer Meteor Soc,* **75,** 811-822.

14 Lanicci, J M, T N Carlson and T T Warner, 1987: Sensitivity of the great plains severe-storm environment to soil-moisture distribution. *Mon Wea Rev,* **115,** 2660-2673.

15 Lean, J and P R Rowntree, 1993: A GCM simulation of the impact of Amazonian deforestation on climate using an improved canopy representation. *Quart J Roy Meteor Soc,* **119,** 509-530.

16 Milly, P C D and K A Dunne, 1994: Sensitivity of the global water cycle to the water-holding capacity of land. *J Climate,* **7,** 506-526.

17 Mintz, Y, 1984: The sensitivity of numerically simulated climates to land-surface boundary conditions. *The global climate,* J Houghton, Ed, CUP, 79-105.

18 Namias, J, 1958: Persistence of mid-tropospheric circulations between adjacent months and seasons. The atmosphere and sea in motion (Rossby memorial volume), *B Bolin, Ed, Rockfeller Institute Press,* 240-248.

19 Nobre, C A, P J Sellers and J Shukla, 1991: Amazonian deforestation and regional climate change. *J Climate,* **4,** 957-988.

20 Noilhan, J and S Planton, 1989: A simple parametrization of land surface processes for meteorological models. *Mon Wea Rev,* **117,** 536-549.

21 Oglesby, R J, 1991: Springtime soil moisture variability, and North American drought as simulated by the NCAR Community Climate Model I. *J Climate,* **4,** 890-897.

22 Rasmusson, E M, 1967: Atmospheric water vapour transport and the water balance of North America, Part I: Characteristics of water vapour field. *Mon Wea Rev,* **95,** 403-426.

23 Rasmusson, E M, 1968: Atmospheric water vapour transport and the water balance of North America: II Large-scale water balance investigations. *Mon Wea Rev,* **96,** 720-734.

24 Rasmusson, E M, 1971: A study of the hydrology of Eastern North America using atmospheric vapour flux data. *Mon Wea Rev,* **99,** 119-135.

25 Rowntree, P R and J A Bolton, 1983: Simulation of the atmospheric response to soil moisture anomalies over Europe. *Quart J Roy Meteor Soc,* **109,** 501-526.

26 Sellers, P, Y Mintz, Y C Sud and A Dalcher, 1986: A simple biosphere model (SiB) for use with general circulation models. *J Atmos Sci,* **43,** 505-531.

27 Viterbo, P and A C M Beljaars, 1995: A new land surface parametrization scheme in the ECMWF model and its validation. Submitted to J Climate.

On the skill and consistency in medium range weather forecasting

Introduction

The interpretation of NWP verification statistics has over the years caused some debate. Improvements in the analysis and forecast systems have not always been reflected in the standard verification scores. While both the Root Mean Square Error (RMSE) and the Anomaly Correlation Coefficient (ACC) for the ECMWF model have improved significantly in the short and early medium range over the last ten years, it has not been the case for forecast ranges beyond day 5 (Fig. 1). The forecast skill also varies from day to day, which is reflected in occasional large changes between successive forecasts, commonly referred to as 'forecast inconsistency' (Fig. 2). Among users of the ECMWF forecasts there is a feeling that improvements in the forecast system should not only lead to reduced forecast errors, but also to improved consistency.

It has become imperative for ECMWF to understand what mechanisms affect the level of RMSE and 'forecast consistency'. This is important because forecasters tend to put more confidence in consistent forecasts than inconsistent ones. It is also important because the absence of any significant reduction of the RMSE level and the slightly increased inconsistency which have been seen over the last few years might be taken as an argument that the level of predictability had been reached.

Improvements of the model

Weather forecast models can be skilful in two ways. The obvious one, generally most appreciated, is their ability to predict the atmospheric features with the right intensity at the right time and in the right place. Most of the work in NWP over the years has focused on getting this better and much progress has been made. But less familiar but equally important is to get the model atmosphere to behave in a realistic way at all forecast ranges. Since the start of NWP development in the 1950's until a few years ago, a constant deficiency of the models, hampering further improvements has been a systematic decrease in the level of eddy kinetic energy in the models for longer forecast ranges. This was seen not only in the statistics of the energetics of the model, but also in the underestimation of developments in the synoptic features on the traditional forecast charts.

A realistic and stable model climate, which creates realistic atmospheric features of all scales, even if they are not predictable *per se*, is necessary to maintain forecast skill. For example, skilful forecast of extensive blockings are often preceded by small-scale intense cyclogeneses. These cannot be forecasted if the fields are too smooth due to low resolution or excessive diffusion.

Already during the first years of ECMWF operational forecasting the short range forecasts of strong blockings and intense storms were fairly good. Up to day 4 they might have been wrong in time and place, but the model was at least able to create such extreme features. Beyond day 4 this was not the case. Instead of blockings the models tended at best to develop smooth ridges; instead of intense lows they created shallow cyclonic pressure systems. This decrease of eddy kinetic energy was due to several factors, most importantly a coarse grid resolution and lack of orographic interaction.

There was, however, one deceptively positive element during this period: *the RMSE scores appeared quite good*. Since intense cyclogeneses and strong blockings are difficult to forecast exactly at the right time and place beyond day 4, any forecast system that avoided or dampened these altogether also escaped the heavy RMSE 'penalties' which followed every time these features were wrongly forecasted. These deficient forecast models thus resulted in unjustified low RMSE levels.

Much of the rationale for higher resolution models during the 1980s was to improve the physical and dynamical realism of the model so that forecasts would be as realistic at day ten as at day one. During the last ten years the ECMWF forecast system has indeed improved both in the definition of the initial state as well as in its general physical and dynamical realism, in particular the elimination of the systematic decrease of eddy kinetic energy. Whilst improvements in the former have contributed to a decrease of the RMSE level, the improvement of the latter has, paradoxically, contributed to an increase, resulting in an almost constant RMSE level over the last several years.

The RMS error level

That a constant level of RMSE does not necessarily imply a constant level of skill can be understood from a decomposition of the RMSE (see Murphy and Epstein, 1989; Simmons, Mureau and Petroliagis, 1994) into three terms which consist of the mean level of forecasted and observed anomalies and the covariance between forecasted and observed anomalies:

$$[1] \qquad E_j^2 = A_a^2 + A_j^2 - 2\,\mathrm{cov}\left((f-c)\,(a-c)\right)$$

where f and a are the forecast and the verifying analyses, c is the climatological value for the verifying time, E_j the RMSE, A_a the observed anomalies, A_j the forecasted anomalies at forecast range j (all quantities averaged over an area and a period of time).

The first term A_a, the atmospheric degree of anomaly or variation around the climate, varies seasonally. The reason why the RMSE is lowest during summer is not because the forecast model is more skilful (it is not) but because the atmospheric eddy kinetic energy activity is at its lowest at that time of the year. This also means that a more extreme winter circulation will yield a higher error level even if the model has been as skilful as during a preceding winter with more normal conditions. Comparisons between years should therefore be made in the same seasons. But even then, a season with higher activity will unavoidably tend to have higher RMSE level.

The second term A_j describes in the same way the model's degree of anomaly or variation around the climate, which of course should be equal to A_a. In the case of an atmospheric model with a systematic decline of eddy kinetic energy the forecasted variation around climate will be lower than the observed one. This will make $A_j^2 < A_a^2$ and contribute to a lower level of RMSE.

The T106 model which was used at ECMWF from 1985 to 1991 was slightly underestimating the atmospheric variability beyond D+4, in particular the formation of strong blockings. The previous T63 model, in operational use from 1983 to 1985, was even more deficient in that way. The fact that the T213 model, with its realistic variability, stays at the same RMSE level as T106, indicates an underlying increase in skill compared to the previous model (the high level of RMSE in the ECMWF model in late 1991 and the first half of 1992 seen in Fig. 1 was due to too high levels of eddy kinetic energy, a temporary problem connected with the numerical scheme formulation).

For long forecast ranges, when the covariance term in [1] tends to 0, the RMSE approaches the value $\sqrt{A_j^2 + A_a^2}$, the error saturation level. A *non-realistic* model which loses eddy kinetic energy and for long lead times approaches a climatological state (where $A_j^2 = 0$) will have 30% *lower* RMSE level than a realistic model (where $A_j^2 < A_a^2$).

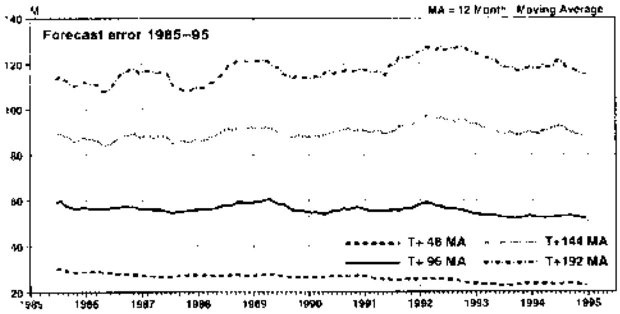

Fig. 1 The ECMWF RMSE scores at 500 hPa over Europe since 1985, day 2, 4, 6 and 8

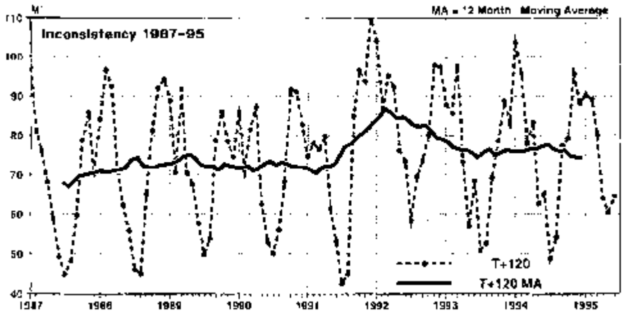

Fig. 2 The inconsistency of the 500 hPa forecast over Europe since 1987, day 5 vs. yesterday's day 6

An example

To illustrate what effect an underestimation of the atmospheric variability has on the RMSE, let us consider a new forecast created by combining the last three ECMWF forecasts with a weighted mean. As can be seen in Fig. 3a, the RMSE of this 'consensus' forecast readily outperforms the operational ECMWF forecast in RMSE terms beyond D+4. Fig. 3b, which shows the level of atmospheric variability, indicates that this decrease in RMSE has been paid by a drastic reduction of realism in the medium range. Fig. 3c shows the average 24 hour activity at 500 hpa in the extratropics in the model compared to the atmosphere.

The level of inconsistency

Since new observations are there to modify the 6 hour forecast used as background to the analysis, forecast changes from one run to the other are unavoidable. Mostly the changes are to the better, even if they are drastic. In a similar way as for RMSE, the inconsistency D_j . measured as the RMS difference between different forecasts valid at the same time, can be decomposed into three terms: the forecasted and observed atmospheric anomalies and the correlation between the forecast anomalies:

$$[2] \qquad D_j^2 = A_f^2 + A_g^2 - 2\,\mathrm{cov}\left((f-c)\,(g-c)\right)$$

where f and g are two forecasts verifying on the same analysis a,. The inconsistency can also be decomposed into the sum of the RMSE of the forecasts and their mutual error correlation *(Persson, 1995)*:

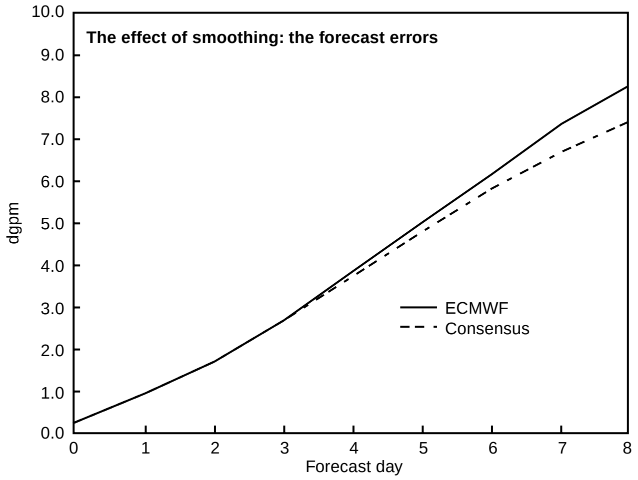

Fig. 3a The RMSE for the northern hemisphere 500 hPa ECMWF forecast averaged from December 1994 to May 1995. The full line denotes the operational forecast, the dashed one a 'consensus' forecast constructed as the weighted mean of the last three ECMWF forecast verifying on the same times. The weights vary with the forecast range range (the shorter the range the higher the weight on the most recent forecast, for longer forecast ranges the most recent forecast is only slightly more weighted than the two others)

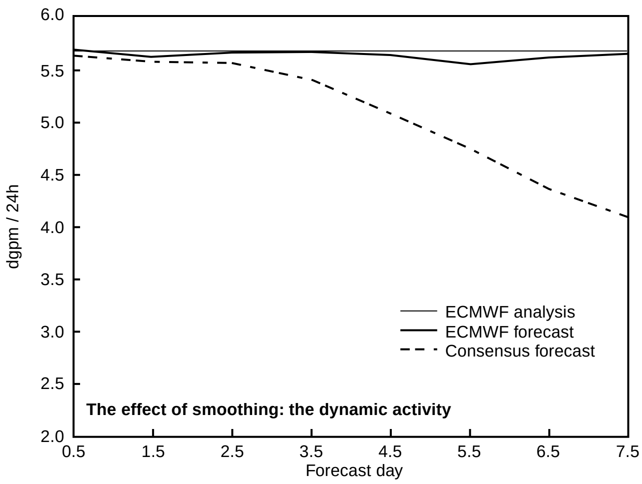

Fig. 3b The level of anomaly over the same period, for the analysis, the T213 forecast and the 'consensus' forecast. The operational model seems to be able to generate anomalous flows with the same frequency as the real atmosphere throughout the 10-day period. The 'consensus' forecast displays a decreasing ability to create and maintain anomalous flow patterns. Models which for numerical or other reasons loose eddy kinetic energy will appear very much as the 'consensus' forecast.

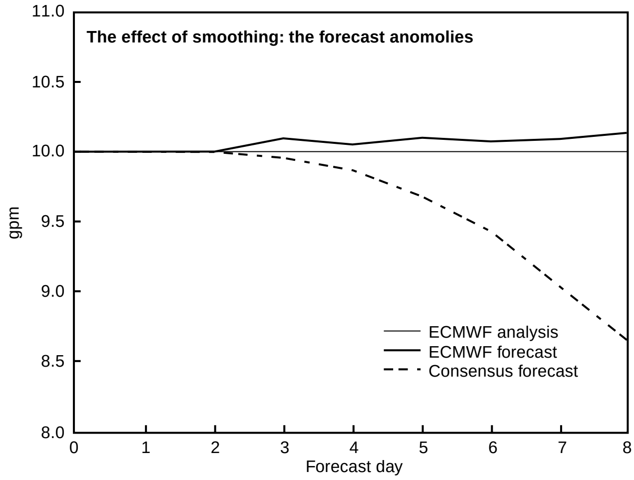

Fig. 3c The average 24 hour 500 hPa tendency over the northern hemisphere extra-tropics over the same period. The diagram shows that lows and highs amplify with realistic intensity even in the later stages of the medium range (when they often are wrongly predicted in time and place) whereas the 'consensus' forecast is less likely to produce strong changes as frequently

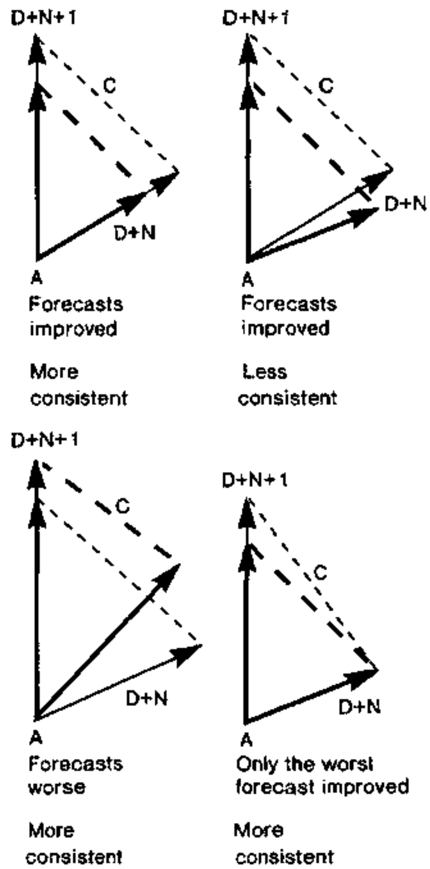

Fig. 4 Vector interpretation of inconsistency. The relation between errors, consistency and dependence between successive forecasts can be nicely illustrated in vector form. The thin arrows indicate the error vectors before a change in the model characteristics has taken place, the thick arrows after the change

[3] $$D_j^2 = E_f^2 + E_g^2 - 2\,\mathrm{cov}\big((f-a)(g-a)\big)$$

where E_f and E_g are the RMSE of f and g over a period of time.

Note that [2] and [3] can refer either to two forecasts both from the same model from different initial times, or to two forecasts from different models at the same time.

The relationship between changes in the performance of the model and its effect on the consistency can be understood with a vector interpretation of [3] (Fig. 4). A reduction of the RMSE of the two forecasts does not necessarily reduce the inconsistency if at the same time the correlation between the forecast errors is reduced. This can happen if more or better observations lead to larger updates of the 6 hour background forecast, for example over the oceans.

Summary

Improvements to a numerical weather prediction system cannot be measured by traditional RMSE and ACC scores only. The lack of significant decrease of the RMSE level of the ECMWF forecasting system over the last ten years has in fact been due to the improvements of the realism of the model over-shadowing the RMSE reducing improvements like better initial conditions. With the T213 model, a model climate close to reality has been reached, and, for the purpose of using the forecast in the medium-range, averaging or smoothing can be quite advantageous for the forecaster. However, any such intervention should be made in the post-processing, not in the actual forecast computation.

The inconsistency is not only linked to the errors of the forecast but also to their correlation. A large decrease should not necessarily be expected from an analysis improvement, since for example an improved observation coverage in a data sparse area will tend to make successive analyses less correlated, and thus contribute to maintain the level of inconsistency. Forecast inconsistencies are difficult to handle for the forecasters who need to avoid conveying abrupt changes of their forecasts to the public. In order to reduce the inconvenience of inconsistency, it is important not to over-interpret the medium range forecasts on scales which are normally not predictable at a certain forecast range. The temptation to over-interpret non-predictable scales is also great when the forecasts are very consistent. For further discussion of the treatment of inconsistencies in the medium range, the reader can refer to the *ECMWF User Guide 2.0*, 1994, pp. 49-54.

Anders Persson, Bernard Strauss

References

1. Murphy, A. and Epstein, E. (1989), Skill scores and correlation coefficients in model verification. *Monthly Weather Review,* **117**, pp. 572-581.
2. Persson, A. (1995): On the statistical relationship between consistency and skill in medium range weather forecasting. *ECMWF, Operations Department Technical memorandum* (under preparation).
3. Simmons, A.J., Mureau, R. and Petroliagis, T. (1995). Error growth and predictability estimates for the ECMWF forcasting system. *Q J R Met. Soc.* **121**, 1739-71.

A new Data Handling System

ECMWF history in mass storage

ECMWF's mass-storage system based on the Los Alamos Common File System (CFS) has been very successful for many years. Given that CFS was designed originally to use Los Alamos' home-grown network hardware and protocols (this was long before TCP/IP), and that it depended exclusively on the IBM 'steam driven' 3850 robotic mass storage unit with 50 MB cartridges, its ability to adapt to subsequent technology and requirements is remarkable. CFS was originally intended to handle of the order of one terabyte of data. At ECMWF today, CFS works reliably in conjunction with TCP/IP over HIPPI channels, and uses StorageTek 4400 silos with 3490-compatible tape units for data storage, handling over 30 terabytes of data, 55000 tape volumes, and traffic in excess of 100GB per day. Availability of the system is of the order of 99%.

Although ECMWF expended a considerable amount of highly-skilled human resources to achieve this state of affairs, it could not have been done without the wholehearted cooperation and active help which we received from Los Alamos staff. The process of attempting to find a worthy successor to it has enhanced our already healthy respect for the system and for the team who built it. Systems of comparable quality are indeed difficult to find.

However, in the present environment at ECMWF, CFS is showing its age. Although the "IBM mainframe plus MVS operating system" combination gives an extremely reliable service, it is an uneasy fit in an environment where the client machines are almost universally running UNIX variants of one sort or another, and in which all the clients are using TCP/IP. There is, for example, effectively no possibility of providing a mounted-file interface (NFS; AFS, DFS) to CFS. Furthermore, the structure of CFS - although highly reliable - does not make it easy to adapt the system to new storage hardware. And, finally, Los Alamos itself has determined that its future data storage needs will not be met with CFS or with a development of it: anything further that is done with CFS at our site must be done with our own expertise and resources alone.

Time does not stand still in other respects. A new supercomputer is planned for ECMWF which will deliver, in 1996, five times the power of the existing Cray C90-16. Two years afterwards, a major upgrade to this system will more than double that power again. Based on past trends (see Fig. 1) this will lead to a substantial increase in the storage capacity being required. Thus the time for a major change to the mass storage system is clearly upon us.

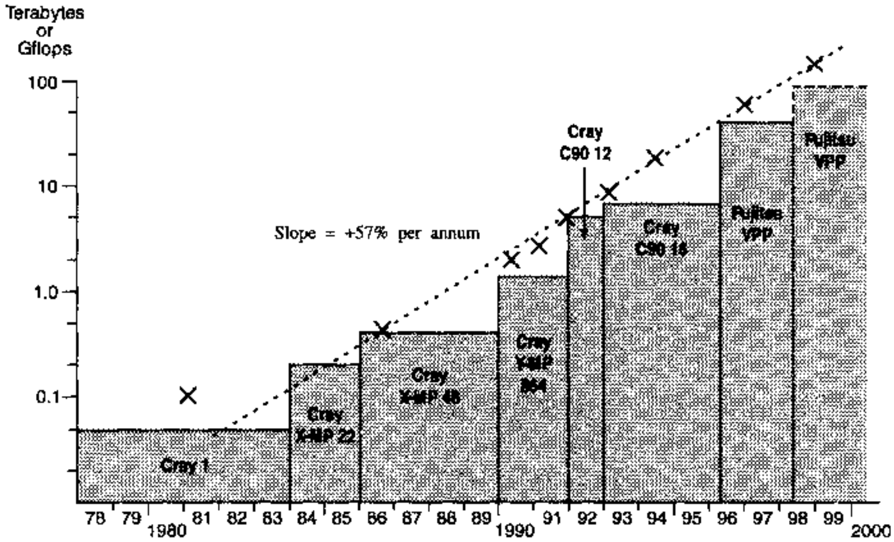

Fig. 1 Showing archive volume in terabytes as X *and computer power in gigflops as histogram bars.*

The next generation system

After much research, thought, and help from other organisations and individuals, ECMWF issued in December 1994 an Invitation To Tender (ITT) for a successor system to CFS. Of the responses received, ECMWF selected the proposal from IBM, based on its RS6000 range of machines and ADSM storage management software.

The new system will be installed in three phases, starting in October 1995. This Phase I Data Handling System (DHS) will be used to provide an ECFILE[1] service to the new supercomputer which will arrive in 1996.

The chosen configuration for Phase I consists of three IBM RS6000 computer systems. Each contains four Power-PC 604 processors and one gigabyte of memory in a symmetric multi-processor configuration. The SPECint(92) rating for the individual processors is 160, hence a total (theoretical) SPECint(92) rating for the installation of 1920.

Each of the computer systems contains four gigabytes of disk installed within the processor cabinets. These disks are intended for use as system residence, boot, dump and swap devices. Each computer system has HIPPI, FDDI and Ethernet interfaces.

The operating system on these machines will be IBM's AIX version 4.1, a POSIX-compliant UNIX system. The operating system uses a journalled file system which provides rapid pick-up of a file system after an unscheduled stop, and the possibility for a different computer to pick up the file system from a failed one. ANSI C, C++, and ANSI FORTRAN 77 and 90 compilers are provided.

There are three disk storage subsystems each containing a total of 125 gigabytes of disk. Part of the available storage space will be configured as mirrored disk pairs, for critical system datasets such as the archive databases; the remaining storage will be configured as RAID (level 3 or 5) for improved availability and higher performance. The disk subsystems are attached to the RS6000 computer systems via IBM's newly announced Serial Storage Architecture (SSA) interface. SSA provides a high-bandwidth (80 megabytes per second per node) connection. It has the additional advantage that the separate storage subsystems and the processors can be interconnected to form a common system, with hardware access from any processor to any disk. This is important, since it forms the basis of the automatic reconfiguration facility if any of the computer systems should fail.

There is a single IBM tape robot model 3494-L10, housed in eight cabinets. One of the cabinets contains the control systems and two IBM 'Magstar' (3590) tape drives, and a second cabinet contains six more tape drives making a total of eight drives in all. The remaining six cabinets store tapes with a total capacity of about 2500 cartridges, or 25 terabytes using Magstar cartridges. An additional two stand-alone tape drives are provided for manually-mounted tapes. Each of the tape drives has dual SCSI interfaces and can therefore be attached to two computer systems simultaneously; though switching between them cannot (in the initial software configuration) be carried out dynamically in response to changing system load.

Software for the system will consist of IBM's ADSM and Turbo ADSM storage manager products, plus the Client Space Manager which moves files between disk and tape according to their level of activity. The IBM HACMP/6000 high-availability monitor will be used to protect the configuration against loss of service. If any hardware or software event occurs which indicates a loss of capability within the system, the HACMP monitor can run reconfiguration scripts which result in the load being taken up by the surviving machines. The network addresses associated with the failed machines can be taken over by the remaining ones, except in the case of the HIPPI connections.

[1] ECFILE is ECMWF's user data storage and retrieval interface and is accessible from all worker systems

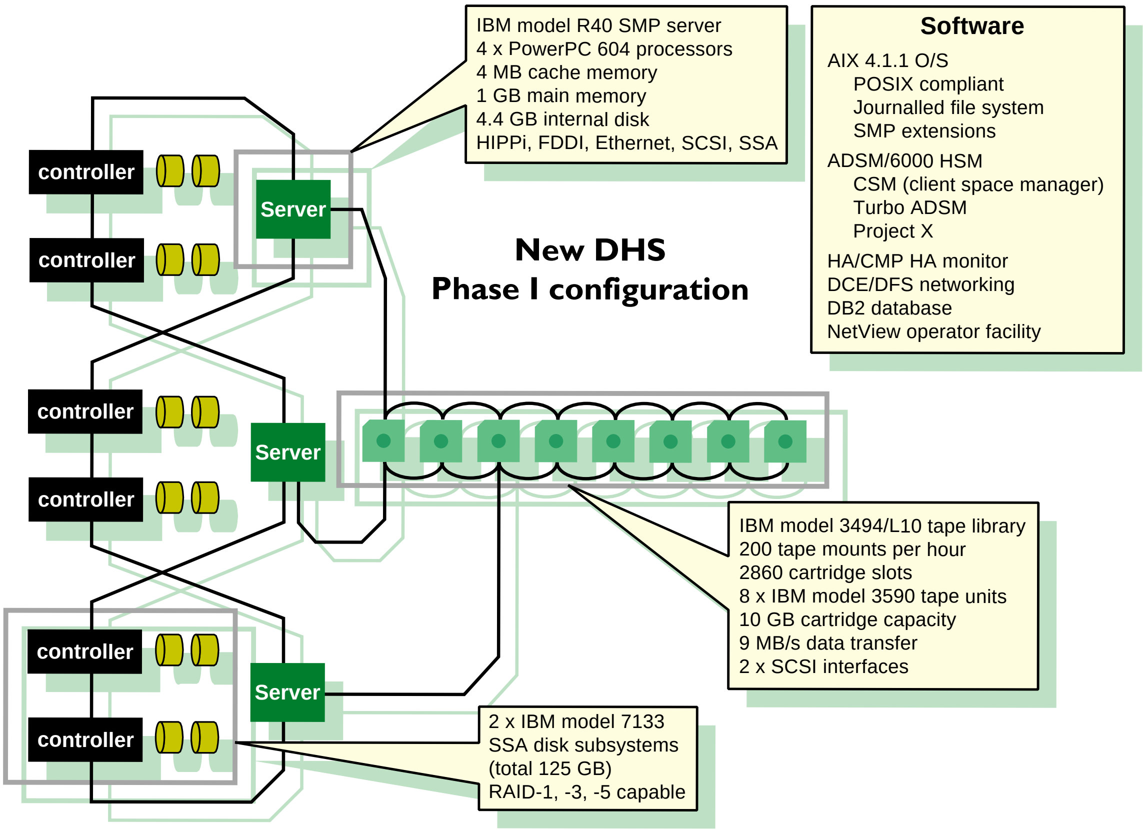

ECMWF will use these three systems in the following fashion:

- One system will be used as a test and development platform which, because of its compatible configuration, can be used as a standby system when not in use for test and development.
- One system will provide the replacement ECFILE service for Centre computer users. The disk space will be configured as cache, with Client Space Manager and ADSM software. User access to this system is expected to be provided via any normal TCP/IP service (NFS, FTP, DFS); direct FTP from clients to ADSM (for large files) will also be provided.
- One system will become the server for the new pilot MARS service. The storage management will be carried out by ADSM, and the MARS server will manage its own disk space in order to optimise the response to user requests.

The MARS and ECFILE servers will be configured via HACMP to take over each other's function in the event of a processor failure; and the development system will be added to the HACMP cluster when not in use for test purpose.

The enhancements presently envisaged for the Phase II and III Data Handling System, which will be installed in 1997 and 1998 respectively, consist of the installation of additional memory on the existing processor configurations, the installation of further processors (one additional cluster at each phase); and finally the installation of further disk capacity and additional disk subsystems to correspond with the increased load. A further 3494 tape robot with eight drives is to be installed for Phase II, and two further identical robots are to be installed for Phase III, corresponding to the growth in archived data volume.

The plan: what ECMWF proposes to do

Our new Mass Storage System will thus be installed in three phases, beginning in October 1995. The first several months of its life will be devoted to configuring and commissioning of the system, and to developing proper operational support procedures for this complex new system. Also during this period we will need to develop support functions for MARS, as it is conceivable that some of the functions needed may adversely impact the stability of the system during testing.

This phase of development is scheduled to end by June 1996, at which time a general-purpose data storage service (the equivalent of our old 'ECFILE') will be provided for the new scientific supercomputer which is being installed in March. Although the configuration available at that time will be only a partial one, the provision of a production service on it offers a useful milestone - an early project verification point. At this stage, ECFILE traffic from other client systems, and all MARS traffic, will still be handled by the old CFS system.

The following year (Spring 1997) in Phase II, the Mass Storage System will be expanded to a configuration which can support up to 100 terabytes of stored data and half a terabyte per day of traffic. This system will take over the production role from CFS in about mid-year (by which time we expect that CFS will have around 60 terabytes of stored data). At this stage, although our operations will need to be able to access data from both CFS and the new system, the entire historic archive of the Centre will still be in the care of CFS.

We will start to copy data out of CFS at this time. We do not count on being able to transfer the old media from CFS to the new system - one reason for changing from CFS is in any case that we expect to have to change to new media. Furthermore, we anticipate that the design of the new MARS system will necessitate some modifications to the file formats that are used; in particular, we wish to increase the file sizes from the present 4-8 megabytes. This means that, as soon as CFS is made read-only in mid-1997, we will immediately start to transfer data from it to the new system. We hope that the full 60 terabytes stored in CFS at that time will not need to be transferred; nevertheless it is clear that at least 2, and possibly up to 5 terabytes per month will need to be transferred over the succeeding 12-month period. During this time we will operate a dual-capability MARS client system which can access both the new and the old servers: if the data is available on the new system it will be accessed from there; otherwise it will be fetched from CFS.

The final phase (Phase III) of the new Mass Storage System installation will start in Spring 1998, coincident with the upgrade of the supercomputer. If all goes well, the data transfer from CFS will be completed about the middle of that year, and the IBM MVS mainframe will be turned off and decommissioned. An era will have come to an end at ECMWF.

Dick Dixon

Secure Member State access to computers using Smart Cards

Over the last few years there has been growing concern about intruders, commonly referred to as ‘hackers’, gaining access to the Centre’s computer services and causing damage in various ways. To combat this threat, the Centre has investigated potential security vulnerabilities, resulting in the definition of a Computer Security Policy together with implementation rules.

One major potential problem identified is the use of cleartext passwords across network links. The Centre realized that something must be done to eliminate the need for specifying those passwords and, after extended discussions with the Member State Security Representatives, presented a proposal to the Centre’s Technical Advisory Committee.

The suggested solution was to introduce user authentication based on ‘smart’ cards. This proposal, an extract of which was published in the ECMWF Newsletter (Autumn 94), was accepted and the SecurID Project was set up. Security Dynamics Inc. was chosen because they were the only company to support a batch environment of the required characteristics, in addition to normal interactive access control.

It also emerged that there would be a need to transfer files from the Centre to the Member States without the use of a password. “Eccopy” provides this facility, as described later.

The launch of the interactive service

To replace the existing password based services (ecserver) a new workstation server, ecgate1, has been set up. Access to ecgate1 is only permitted via a valid SecurID card code.

SecurID Smart Cards have been sent out to the Member State Computer Representatives, who distribute and administrate the cards within their Member State.

Officially the service on ecgate1 became available at the end of May. The service on ecserver (munin) will be phased out on 31 October 1995.

At that time direct connections to UNICOS (C90/T3D) will no longer be possible from Member State hosts. To connect interactively to UNICOS, Member State users will have to login on ecgate1 with their SecurID Card and then connect to UNICOS. However, direct FTP access to UNICOS will still be permitted for some time.

Eccopy

Eccopy has also been made available to the Member States. This allows a Member State user who has given the relevant permission on the UNIX workstation he/she is using at the Member State location, to copy a file from an ECMWF computer to a designated local directory.

This facility will also be used to transfer back to the Member State any output from batch jobs authenticated with SecurID Cards.

Other services

Already now it is possible to submit local batch jobs on ecgate1. That includes MARS jobs.

The Centre will provide facilities to submit batch jobs soon and MARS requests later on direct from Member State UNIX workstations, again based on SecurID Card authentication. The batch job submission can either use NQS or special client software.

These facilities use the concept of a “certificate”, which can then be sent along with any request. Such “certificates” will be created by a special utility after prompting the user for a card code. They will be re-usable for batch job management, for MARS requests, and in exceptional circumstances for batch job submission. They will be valid only for a fixed period. The necessary precautions will be taken so that only the requesting user can use the certificate.

Conclusion and outlook

The additional security provided with SecurID Card authentication has led us to investigate the possibility of allowing access to the Centre not only via the dedicated Member State links but also via our general Internet link. If such a service can be provided in a secure fashion, then ECMWF staff members on mission and Member State users may be permitted to access the Centre directly via the Research Networks.

Dieter Niebel

ECMWF Publications

Technical Memoranda

No. 212: The sensitivity of winter evaporation to the formulation of aerodynamic resistance in the ECMWF model. Beljaars, A.C.M. and P. Viterbo. February 1995.

No. 215: The parametrization of surface fluxes in large scale models under free convection. Beljaars, A.C.M. February 1995.

No. 216: Variational retrieval of humidity profile, wind speed and cloud liquid water path with SSM/I: potential for numerical weather prediction. Phalippou, L. July 1995

No. 217: Impact of the ERS-1 scatterometer wind data on the ECMWF 3-D VAR Assimilation system. Gaffard, C. and H. Roquet. July 1995.

No. 218: A new sub-grid scale orographic drag parametrization: Its formulation and testing. Lott, F. and M. Miller. July 1995.

No. 219: Predictability studies with high resolution singular vectors. Buizza, R., R. Gelaro, F. Molteni and T.N. Palmer. August 1995.

No. 220: Estimating the covariance matrices of analysis and forecast error in variational data assimilation. Fisher, M. and P. Courtier. August 1995.

Technical Reports

No. 75: An improved land surface parametrization scheme in the ECMWF model and its validation. Viterbo, P. and A.C.M. Beljaars. March 1995.

No. 76: Global Sea Ice Concentration Data Set for use with the ECMWF Re-Analysis System. Nomura, A. March 1995.

Forecast and Verification Charts

to July 1995.

Seminar Proceedings

Seminar Proceedings of sub-grid scale physical processes. 5-9 September 1994.