In this issue

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, RG2 9AX, UK

Fax: +44 118 986 9450

Telephone: National 0118 949 9000

International . . . . +44 118 949 9000

Internet URL . . . . http://www.ecmwf.int

Cover

Clouds and radiation have a large impact upon sea-surface temperatures - see the article on page 2 for some recent changes in this area.

Editorial

Some significant changes were introduced into the operational model in December 1997, including changes to the radiation, convection, and cloud parametrizations. An overview of these changes is presented on pages 2 including a discussion of their impact on both seasonal simulations and 10 day forecasts.

ECMWF's Ensemble Prediction System (EPS) is being more and more widely used. Various ways have been sought to reduce the vast amount of data generated to more manageable proportions. Clustering has been one way to achieve this up to now, however 'Tubing'could be a viable alternative. The article on pages 7 explores this idea further.

Four more of ECMWF's Technical Memoranda are summarised on pages 14 to 16.

The success, or otherwise, of optimising large codes on any big computer system depends a lot on the tools available that can analyze and report on the behaviour of that code as it executes. A new tool (SALTT) has recently become available, the article on page 16 discusses its first application to a large code, namely ECMWF's Integrated Forecast System (IFS).

Finally, a warning to all users of ECMWF's computer service: the old file storage service (ECFILE) will close on 31 July 1998 – see page 20.

Changes to the Operational Forecasting System

Recent changes

The method used for the computation of initial perturbations of the Ensemble Prediction System (EPS) has been changed on 24 March 1998. From this date, initial perturbations are a combination for day D of fastest growing modes between D-2 and D ('Evolved Singular Vectors') and fastest growing modes between D and D+2; the meteorological impact of this change is minor (slightly increased spread in the early medium range).

The reference model spectral resolution was increased from T213 to T_L319 on 1 April 1998 (improved orography, reduced diffusion, but Gaussian grid unchanged). The other modifications introduced on this occasion (Cy18r5) were:

1. a new 2-time level semi-Lagrangian numerical scheme; this made possible the use of a Linear Gaussian Grid reflected by the T_L notation already introduced for the EPS configuration in Dec. 1996;

2. use of a new 2'30" orography reference dataset; as a result, some large errors over Antartica have been corrected, and spurious noise in other areas suppressed;
3. improved sea-ice limits determination.

Planned changes

- Coupling of the atmosphere and ocean-wave model for data assimilation and forecasting;
- Use of both significant- and standard-level winds, temperatures and humidities from radiosondes (geopotential will no longer be used);
- Use of extra off-time data, mostly SYNOPs and DRIBUs (hourly data will be accepted, which is now only the case for mobile platforms such as aircraft);
- Use of 1D-var SSM/I total column water-vapour;
- Extension of the use of GOES winds to the northern extratropics.

François Lalaurette

Introduction of revised parametrizations of physical processes into the IFS

Together with changes to the data assimilation system and dynamical component of the model, over the last 10 years several changes to the parametrizations of the ECMWF model have brought about improved model performance (such as the revised sub-grid scale orography and prognostic cloud schemes introduced in April 1995 - see Summer 1995 Newsletter No. 70). On the 16th of December 1997 a new set of revised parametrizations was introduced into the operational version of the IFS (as CY18R3), with changes being made to the radiation, convection and cloud parametrizations. A change was also made to the numerical treatment of vertical diffusion.

The revisions were undertaken for several reasons. Firstly to improve the physical basis of the parametrizations and their performance as measured against observations and detailed models (such as line-by-line radiation codes and fine-scale cloud resolving models for convection). Secondly the changes were aimed at correcting errors in the Top of Atmosphere (TOA) and surface energy budget, important for coupled ocean-atmosphere simulations of the model used in experimental seasonal forecasting. Some improvement is also seen in aspects of the 10-day forecast performance of the model.

In this article the changes will be briefly described and their impact upon both seasonal simulation and 10 day forecasts discussed.

Description of changes

Radiation

The treatment of the water vapour continuum in the long-wave part of the radiation code was updated to better match line-by-line computations, correcting an overestimation of clear-sky cooling in the lower troposphere and an underestimation of cooling in the upper troposphere. The short wave radiation scheme was made more flexible, being able to be used with either 2 or 4 spectral bands (operationally 2 bands are used). Short wave cloud optical properties were revised to remove excessive in cloud absorption and a temperature dependent effective radius for ice particles introduced. A treatment of cloud inhomogeneity following *Tiedtke (1996)* was included, the liquid water path used by the radiation scheme being multiplied by 0.7.

Convection

The Tiedtke (1989) mass flux convection scheme has been used in the operational model since 1989. It allows deep convection to form if moisture convergence into a column of the atmosphere is positive. However recent studies have indicated that a direct link between moisture convergence and convective activity may lead to poor simulation of synoptic variability in the tropics *(Slingo et al., 1994)*. Experience also suggests that the model precipitation pattern in the tropics may be less broad (in the Intertropical Convergence Zone (ITCZ) - for example) than observed. In the revised scheme the presence of deep convection is determined by the depth of instability; if the cloud depth exceeds 200 hPa it is deemed to be deep convection, shallow convection otherwise. Also for **deep convection** the estimation of cloud base mass flux is changed from one using an assumption of the quasi-equilibrium of the moisture content of the sub-cloud layer when convection is active, to assuming that convection acts to reduce Convective Available Potential Energy (CAPE) towards zero over a specified time-scale *(Nordeng, 1994)*.

Cloud scheme

The prognostic cloud scheme *(Tiedtke, 1993)* was introduced into the operational model in 1995 (see Summer 1995 Newsletter No. 70). In that original version ice falling out of a layer is converted to snow and assumed to fall to the surface, undergoing evaporation, within a single timestep. The revised scheme has an improved mathematical treatment, allowing ice to fall from one level to the next within a cloud, although the original formulation is retained for ice falling into clear sky. This treatment was tested at the time of operational implementation of the Tiedtke cloud scheme, but the increased cloud fractions and water contents in the upper troposphere were found to cause a spurious warming in association with excessive in-cloud absorption of short-wave radiation in the radiation scheme at the time. As noted above this has now been removed through a revision of cloud optical properties.

Vertical diffusion

This is now iterated three times within a model time-step, resulting in a more accurate estimate of the surface drag coefficient. This is beneficial when the ocean wave model is running coupled to the atmospheric component of the IFS, but in the stand alone tests described here has virtually no impact.

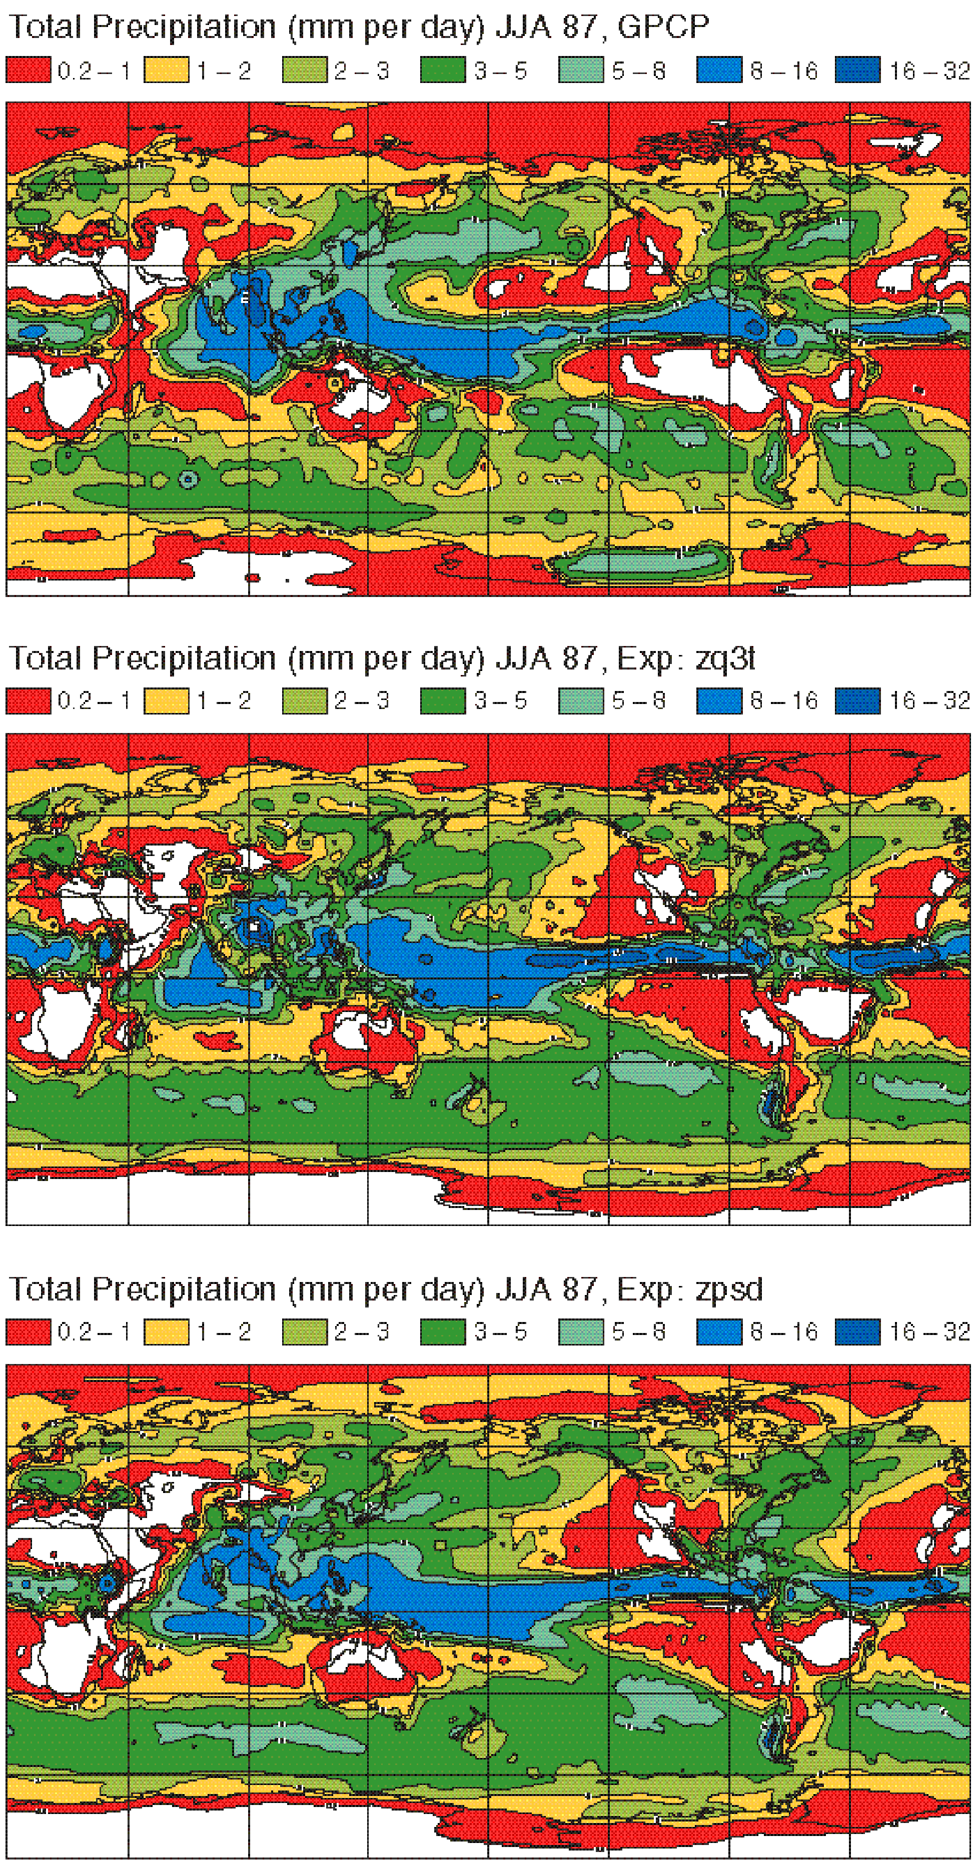

Figure 1: *Comparison of (top) GPCP precipitation climatology for June/July/August 1987 with average precipitation over an ensemble of three seasonal forecasts with the IFS at T63 for the same period using (middle) the control physics package and (bottom) the revised physics package.*

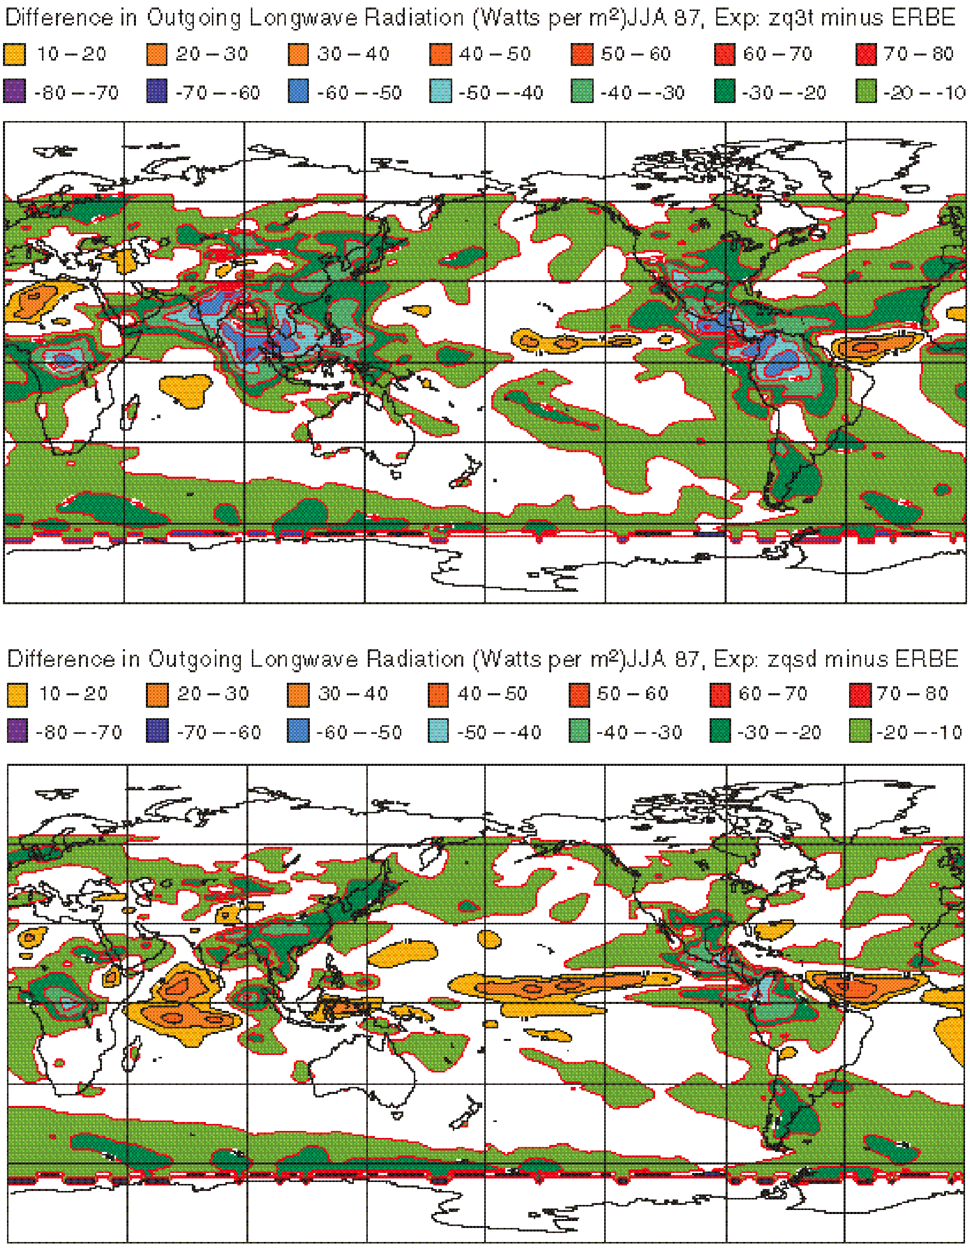

Figure 2: *Difference between simulated Top of Atmosphere Outgoing Long-wave Radiation (OLR) from ERBE data for June/July/August 1987 averaged over an ensemble of three seasonal forecasts with the IFS at T63 using (a) the control physics package and (b) the revised physics package.*

Impact on seasonal simulations

A coupled ocean-atmosphere version of the IFS at T63 resolution is now used routinely at ECMWF as part of its seasonal forecasting activity. Of importance here is the prediction of tropical sea surface temperature (SST) anomalies, which are affected by both surface fluxes of heat and moisture, and also the surface wind stress which is related to the horizontal distribution of heating by physical processes in the atmosphere, especially convection.

The introduction of the revised parametrizations results in an improved tropical precipitation distribution, mainly as a result of the change to the convection scheme. Figure 1 shows mean surface precipitation from an ensemble of three T63 simulations (started one day apart and using observed SST) averaged over June/July/August 1987 compared with climatology (GPCP). Over India rainfall is more evenly distributed rather than being concentrated into certain locations. The intensity of the ITCZ north of the equator in the Pacific and Atlantic is reduced. Both these changes bring the simulation into better agreement with climatology. Associated with these changes is a reduction in the strength of the Hadley circulation and also trade winds over the Pacific and Atlantic.

Radiative fluxes at the top the atmosphere and at the surface are also improved. Figure 2 shows differences between simulated Outgoing Long-wave Radiation (OLR) and satellite measurements for June/July/August 1987. Large decreases in negative errors (implying too high an OLR in the model) are seen in the tropics, due to a combination of the revised radiation scheme plus changes to the cloud distribution and amount brought about by the convection and ice fallout changes. In mid-latitudes errors are also reduced by the revision of the long-wave radiation scheme.

The change to the ice fall out formulation causes both upper level cloud amounts and ice water content to increase. Although the ice content of the atmosphere is

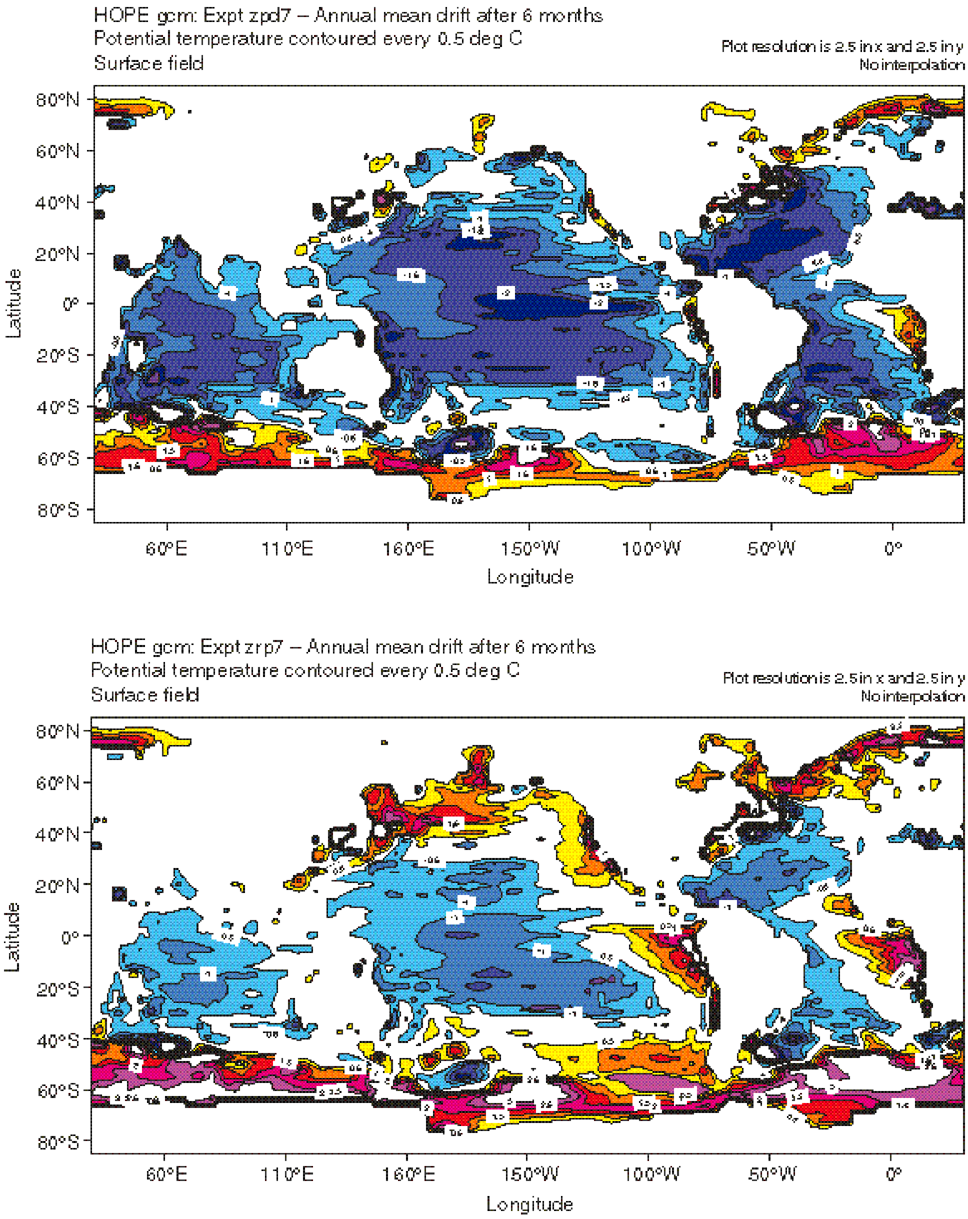

***Figure 3:** Drift of sea surface temperature (as compared to observations) for the last month of a six month forecast at T63 with the coupled ocean-atmosphere version of the IFS, averaged over an ensemble of 24 forecasts (starting from the 1st January, 1st April, 1st July and 1st October for years 1991 to 1996) using (a) the control physics package and (b) the revised physics package.*

not well observed, previous studies (Rizzi and Jakob, 1996) have suggested that the original ice fallout formulation leads to an underestimate of the ice water amounts. Study of a well documented FIRE case *(Klein and Morcrette, 1997)* indicates that the increased ice water contents are in better agreement with observed amounts.

Similar changes to those described above also occur in winter 1987/88 simulations. In this season the changed distribution in heating due to convection and clouds in the tropics leads to stronger westerly flow in the upper troposphere of the central and east Pacific, in better agreement with observations. The upper level westerly flow in this region is important for tropical-extratropical interactions which in turn play a role in determining blocking over the North Atlantic and Europe.

Overall the changes in the precipitation distribution and radiative fluxes lead to an improved surface energy balance, with generally more heat going into the ocean. Initial coupled ocean-atmosphere simulations indicate that the SST bias over a large part of the tropics is reduced by the revised parametrizations. Figure 3 shows the annual mean drift in SST after 6 months for an ensemble of forecasts obtained using the coupled ocean-atmosphere version of the IFS with the new and old physics. Errors have been substantially reduced in the tropical Indian, Pacific and Atlantic oceans, although warm biases have increased in the vicinity of stratocumulus sheets and in the southern hemisphere depression tracks due to remaining deficiencies in cloud amounts in these regions. Initial indications are that the ability of the revised model to predict seasonal sea surface temperature anomalies appears to be slightly increased. The changes to the climatology of the model will also impact upon the performance of the model in a 19-year simulation (1979-1997) to be carried out shortly as part of the AMIP-II project.

Impact on 10-day forecasts

Several series of experiments were carried out during testing of the package of parametrization changes to assess their impact on data assimilation and 10-day forecast skill. In 4D-Var experiments the impact of the package is neutral, although during pre-operational testing in December there was a small improvement in skill in both northern and southern hemisphere scores.

As in the seasonal simulations, the change to convection has an impact upon the distribution of precipitation in T213 forecasts. The ITCZ in the Pacific and Atlantic does not intensify as much over the first 5 days of a forecast than with the operational scheme, giving better agreement with observed estimates. Also looking at precipitation amounts from individual forecasts it is noticeable that the new scheme gives smoother rainfall patterns over both the tropical and mid-latitude oceans.

The changes have a non-negligible impact upon the temperature bias of the model. Figure 4 shows zonal mean temperature differences between 5-day forecasts and analysis for a series of 15 forecasts in September 1997 carried out using an experimental version of the 4D-Var assimilation system (introduced into operations in early December 1997). The new parametrizations substantially reduces the cold bias of the control forecasts, especially in the lower troposphere of the tropics and sub-tropics. This is mainly due to the revised long-wave radiation scheme. In the upper troposphere the warm bias at 200 hPa in the tropics is removed, while in the mid-latitudes the cold bias in the upper troposphere is increased.

As mentioned previously, although the new radiation scheme is a better fit to detailed radiation calculations, it cools the upper troposphere more that the previous operational code. Alone it produces more substantial increases in the cold bias of the model in the upper troposphere which are detrimental to forecast skill. Only by combining the revisions to the radiation scheme with the ice fallout and convection changes, which tend to increase upper level cloud cover and warm these regions, is acceptable forecast skill maintained. This illustrates an important point when developing new parametrizations. Although a revised parametrization may provide a better physical representation of a process, comparing well with observations and detailed models, improved model performance is not guaranteed. Deficiencies in other parts of the models parametrization package may contribute to a worsening of performance, and interactions between different components of the models parametrization need to be taken into account.

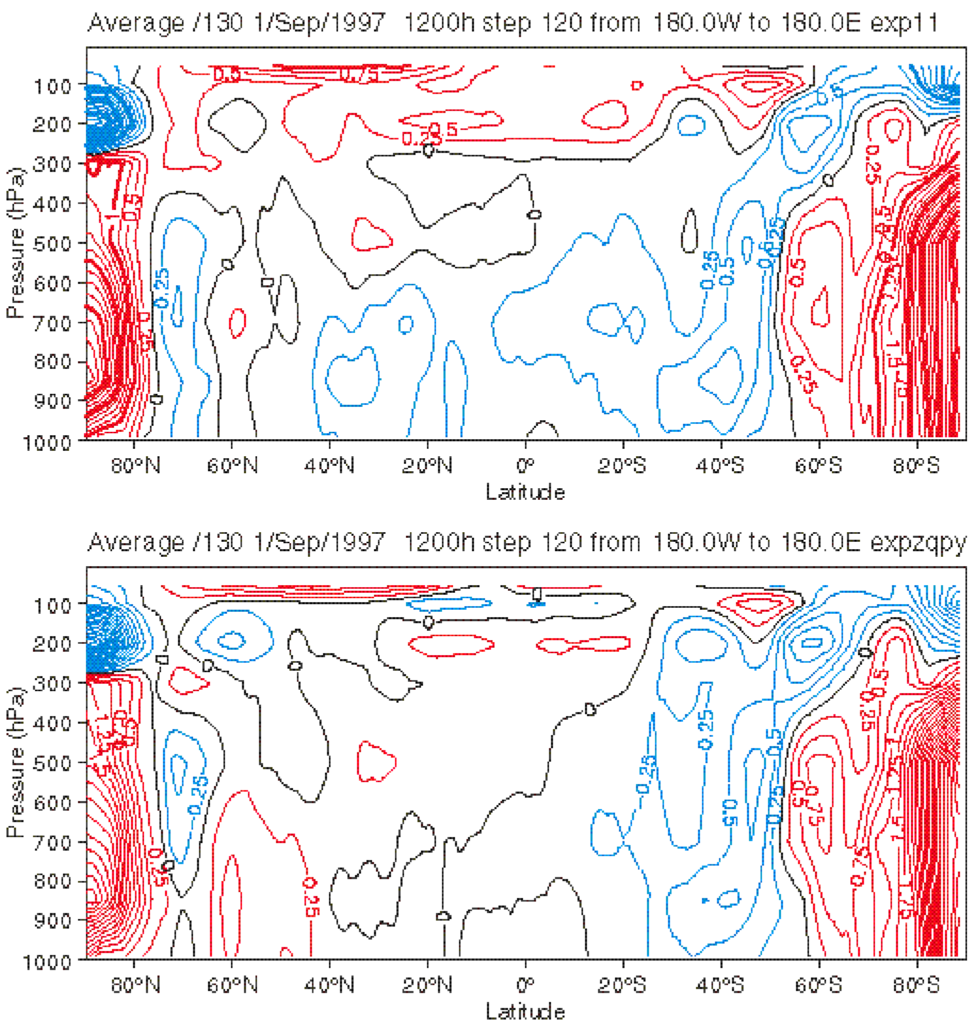

Figure 4: *Zonal mean temperature error at day 5 from a series of 15 forecasts for September 1997 using (a) the control physics package and (b) the revised physics package.*

Summary

A revised set of physical parametrizations for the IFS has been described. Apart from a general reduction of model temperature bias, the direct impact upon medium-range forecasts is small. However, there is a large impact upon the ability of the model to accurately capture many aspects of the tropical climate on seasonal timescales. The improved distribution of heating and TOAand surface energy budgets is of benefit to coupled ocean-atmosphere modelling at ECMWF and hence to forecasting on seasonal time-scales.

Acknowledgements

The help of Tim Stockdale and others in the seasonal forecasting group at ECMWF in the production of figure 3 is acknowledged. ❑

References

Klein, S. and **Morcrette, J.-J.**, 1997 : Simulation of a cirrus cloud observed during the FIRE-II field experiment., ECMWF Research Department Memorandum, ECMWF, Shinfield Park, Reading, Berks, United Kingdom, August 1997

Nordeng, T.E., 1994 : Extended versions of the convection parametrization scheme at ECMWF and their impacts upon the mean climate and transient activity of the model in the tropics., Research Department Technical Memorandum No. 206, ECMWF, Shinfield Park, Reading, Berks, United Kingdom

Rizzi, R. and **Jakob, C.**, 1996 : Evaluation of model OLR in cloudy regions using TOVS 1b data., ECMWF Research Department Memorandum, ECMWF, Shinfield Park, Reading, Berks, United Kingdom, December 1996

Slingo, J.M., **Blackburn, M.**, **Betts, A.**, **Brugge, R.**, **Hoskins, B.J.**, **Miller, M.J.**, **Steenman-Clark, L.** and **Thurburn, J.** (1994) Mean climate and transience in the tropics of the UGAMP GCM: Sensitivity to convective parametrization, *QJR. Meteorol. Soc.*,**120**, 881-922

Tiedtke, M., 1989 : A comprehensive mass flux convection scheme for cumulus parametrization in large-scale model., *Mon. Wea. Rev.*, **117**, 1779-1800

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

Tiedtke, M., 1996 : An extension of cloud-radiation parametrization in the ECMWF model: The representation of sub-grid scale variations of optical depth., *Mon. Wea. Rev.*, **124**, 745-750

David Gregory, Jean-Jacques Morcrette, Christian Jakob and Anton Beljaars; Physical Aspects Section

Tubing: an alternative to clustering for EPS classification

Ensemble prediction generates huge amounts of data. The main problem for operational use is the design of manageable products for potential users. The daily output of the ECMWF operational Ensemble Prediction System (EPS) is a collection of forecasts, which are different possible realisations of the meteorological future. From this collection, any weather parameter probability distribution can be derived. This set of distributions, provided they are reliable, should fulfil all operational requirements, since the probability of virtually any meteorological event can be extracted from them.

On the other hand, end users are not necessarily satisfied solely with probabilities. They certainly appreciate *explanations*, designed to help them take their decisions from the probabilistic information. These explanations are not supposed to add any information to the probabilistic forecast, but to allow the end user to understand why the probabilities of weather elements occur, by mentioning the various meteorological possibilities. This 'added value' can only be the result of an *interpretation* of ensemble products by an experienced forecaster, able to understand for instance why a 20% probability of strong precipitation is associated with a given distribution of ensemble forecasts.

The problem is that the distribution of ensemble forecasts is rather difficult to interpret on a meteorological basis: it is almost impossible – in operational conditions – to catch the differences and similarities between 51 forecasts, even over a limited area. Different products have been designed in order to facilitate the forecasters interpretation. Some allow visualisation of specific aspects of the distribution, like the ensemble mean field used as a smooth forecast of the large scale pattern, or the ensemble spread indicating the uncertainty associated to the different meteorological features. Other products are based on an objective classification of ensemble members, indicating the main alternative meteorological patterns. The tubing method belongs to that last category and is proposed as an alternative to conventional clustering techniques.

The tubing algorithm

The idea of tubing is based on the following assumptions:

- Ensemble forecast distribution can generally be considered as monomodal.
- The verification is most likely to be found in the vicinity of the ensemble mean.

- Ensemble modes are generally not reliable: in case of multi-modality the verification is still more likely to be found near the ensemble mean than in the vicinity of ensemble modes.

The tubing is a method of classification of ensemble forecasts that suits monomodal distributions, gives large emphasis to the ensemble mean, and indicates a highly contrasted picture of the possible alternative scenarios.

Unlike conventional clustering methods, the tubing algorithm does not group the members around hypothetical centroids, but along axes coming from the ensemble mean and reaching the extremes of the distribution. If the ensemble distribution is reliable, these axes should represent the directions along which the verification is likely to deviate from the ensemble mean. The algorithm is illustrated figure 1:

- The *central cluster* is obtained by grouping a certain number of those members lying around the ensemble mean. The number of members depends on the tubing configuration, detailed in the next section.
- The member which is the most remote from the ensemble mean is located. It becomes the first *extreme* and defines a *tube* grouping those members lying in the cylinder whose axis of symmetry goes from the ensemble mean to the extreme of the tube, and whose radius is the same as for the central cluster.
- The process is then iterated until all members are classified, with the following rule: a member belonging to a tube cannot become an extreme (but can still be classified in another tube).

The central cluster mean field, similar to the ensemble mean field by construction, indicates the main meteorological option to follow in a deterministic approach, the large scale forecast which is the most likely to verify. The tubes are not averaged but represented by their extremes: since the tube members are not grouped around a centroid but along an axis, a tube mean does not make much sense. However the main information is not to be found in the extremes of the tubes themselves, but in the meteorological differences between these extremes and the central cluster, which indicate the possible *deviations* from the main option.

The tubing configurations

The main parameter of the tubing algorithm is obviously the condition limiting the population of the central cluster. Two main configurations can be distinguished, whether this limitation depends on the actual ensemble spread (spread dependent configuration) or on the expected, 'climatological' ensemble spread (seasonal dependent configuration).

In a spread dependent configuration, eg. central cluster variance limited to 50% of the total ensemble variance, the population of the central cluster varies little from day to day. The central cluster internal spread follows the variations of the ensemble spread, and so does the smoothness of the averaged field, high uncertainty resulting in a smoother field. On the other hand the number of tubes varies little from day to day.

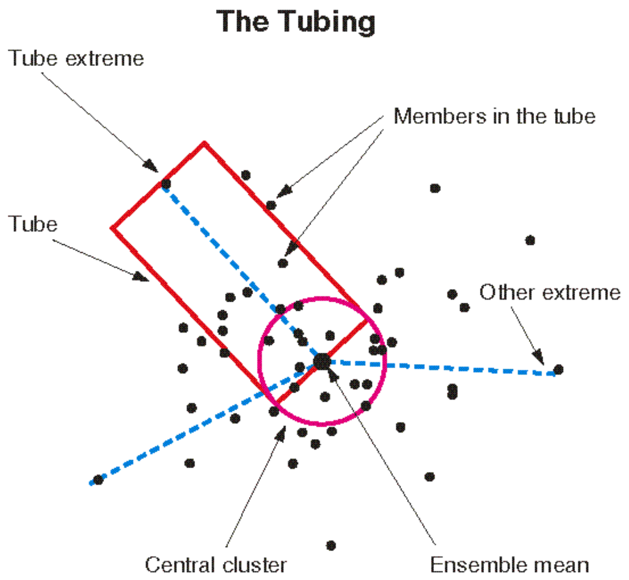

Figure 1: *Schematic representation of the tubing algorithm. The central cluster groups those members lying around the ensemble mean. The tubes are defined by their extremes, which are the most remote members from the ensemble mean.*

In a seasonal dependent configuration, on the contrary, the smoothness of the central cluster mean field varies little from day to day, since the internal spread follows the seasonal, slow variations of the meteorological uncertainty. On the other hand the number of tubes varies from day to day according to the actual ensemble spread, ie. the actual meteorological uncertainty: there may be no tube at all, when all members are grouped in the central cluster, or plenty of them when the spread is especially large.

Both configurations apply similarly to agglomerative clustering algorithms. Aseasonal dependent limitation of the clusters standard deviation is applied for instance to the ECMWF operational clustering. This configuration appears indeed the most suitable for operational purposes, since it allows forecasters to directly link the number of possible meteorological variants – given by the tubes or the clusters – to the forecast uncertainty. The spread dependent configuration is rather suitable for verification purposes and case studies, since it allows one to visualize a given proportion of information (ie. of variance) from the ensemble distribution.

Tubing and clustering

The differences between clustering and tubing are fundamental: clustering algorithms focus on ensemble forecasts similarities, whereas tubing groups ensemble forecasts according to the fact that they *differ* similarly from the ensemble mean. To illustrate their differences the tubing and the operational ECMWF clustering algorithm have been applied to a 2-dimensional ensemble of 51 points randomly extracted from a Gaussian distribution. This idealised ensemble, although roughly monomodal, exhibits local sampling accumulations (figures 2 and 3). A similar

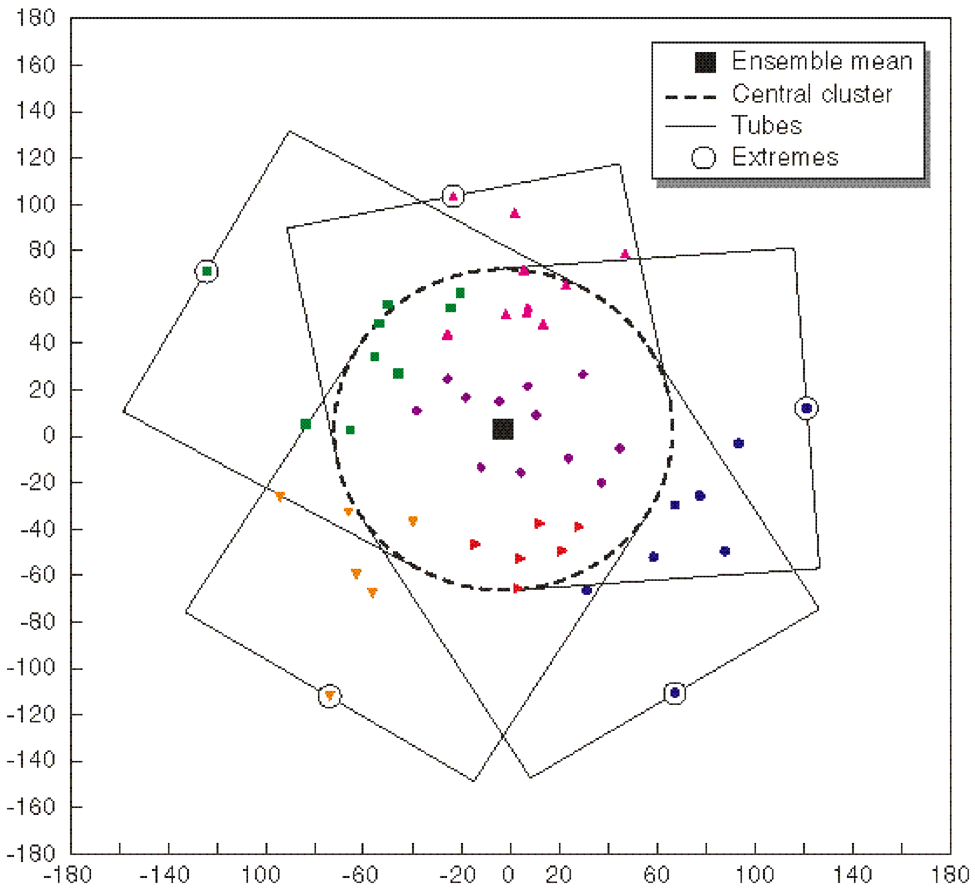

***Figure 2:** Tubing applied to a 2-dimensional idealised ensemble extracted from a Gaussian distribution, giving a central cluster grouping 32 members and 5 tubes.*

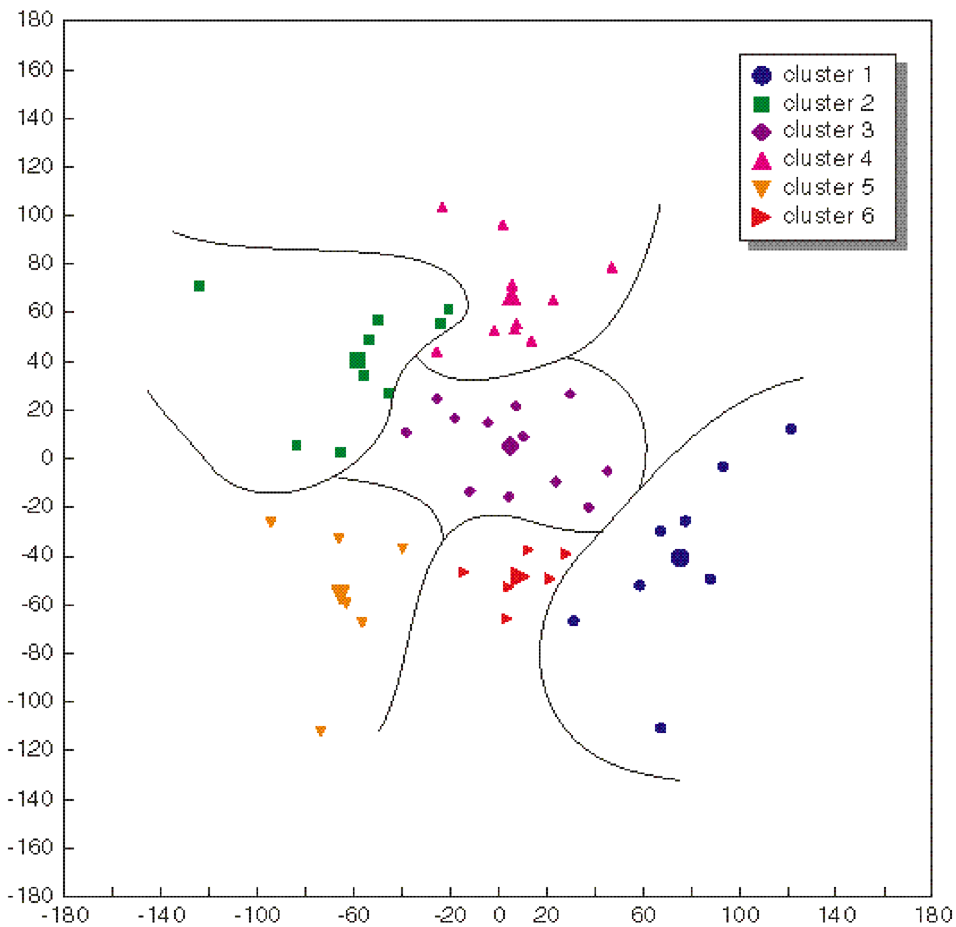

***Figure 3:** Same as figure 2, but clustering algorithm. The clustering classification gives 6 clusters, grouping 8, 9, 12, 10, 6 and 6 members.*

configuration has been used to limit the population of the tubing central cluster and the Ward clusters (st. dev.<50 m).

The Ward algorithm is clearly not suitable for such a distribution: the boundaries between the clusters are arbitrary, often separating members very close to each other; the cluster centroids, close to each other and to the ensemble mean, give a poor representation of the ensemble distribution. On the contrary the tubing classification widely ‘covers’ the whole distribution by localising its centre and its extremes. The tubes boundaries appear as arbitrary as the clusters boundaries, but this should not be considered as a drawback since the tubes are not supposed to be averaged. The tube extremes indicate possible departures from the ensemble mean. Because they are extracted from the low density part of the distribution, tube extremes are very sample dependent compared with clustering centroids. This dependency, which should always be kept in mind when using the tubing product, is the price to pay in order to get more significant indications about the distribution of ensemble forecasts.

Meteorological interpretation based on the tubing classification

The usefulness of the tubing product is discussed in this section on a forecast case study. The +144h EPS forecast based on 22/01/1997 had a large spread (84 m) over Western Europe. Aridge was forecast by the T213 model (as the EPS control, not represented since almost identical) to build up off Ireland and a consequent Northerly anticyclonic flow to establish over Britain and France (figure 4). The ensemble spread mainly related to the development of this ridge and its effects on the circulation over Western and Central Europe.

Tubing

The tubing classification has been applied to this forecast with a 60 m threshold for the central cluster standard deviation (figure 5). The central cluster clearly supports the control forecast: with a 60 m threshold the averaged field is hardly smoothed, and the main synoptic features are still discernable. Four tubes indicate different possible deviations from the main option:

Tube 1 The Northerly flow might be more cyclonic and stronger over the British Isles.

Tube 2 On the contrary the ridge might develop further into the continent.

Tube 3 The ridge might evolve into a dipole with a cut-off low over SW Europe.

Tube 4 A large scale cut-off low might affect Central Europe.

The two first tubes indicate the main, larger possible deviations. At the opposite extreme the last tube is not very different from the control forecast. The 3rd tube proves to indicate the ‘right’ deviation from the ensemble mean: the verification pattern does not match the extreme of the tube itself, but can be subjectively localized somewhere between the ensemble mean and the extreme of the tube (see verification figure 5, bottom right).

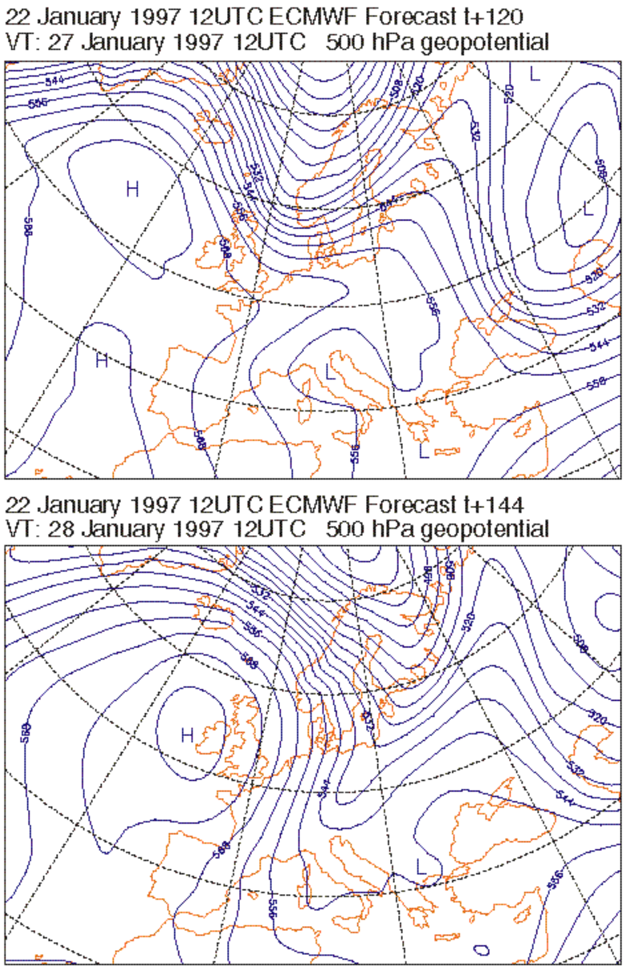

Figure 4: *T213 ECMWF model +120h/+144h forecast based on 22nd January 1997, 500 hPa height.*

Figure 6 shows all the members grouped along the first tube axis, together with the control and the closest member to the ensemble mean: from the closest member to the central cluster boundaries, located 104 m from the ensemble mean, exhibiting a rather anticyclonic North-Westerly flow, through gradually more cyclonic patterns, up to the extreme at 190 m. It is noticeable that some members, eg. the 5th one, do not fully support the tube tendency given by the extreme. As in any classification, the loss of information can only be compensated by an increase in resolution, ie. an increase of the number of tubes, obtained by limiting the size of the central cluster. But this would be at the expense of the conciseness that is of primary importance for operational use in a forecasting office.

Clustering

AWard clustering has been applied to the same situation, giving 5 clusters with a similar configuration (figure 7). The second, most populated cluster, exhibits more or less the same pattern as the control or the tubing central cluster, whereas the 1st and the 3rd ones seem to support the 2 first tubes options, although not as strikingly as the tubes extremes. The 4th and the 5th clusters support more clearly

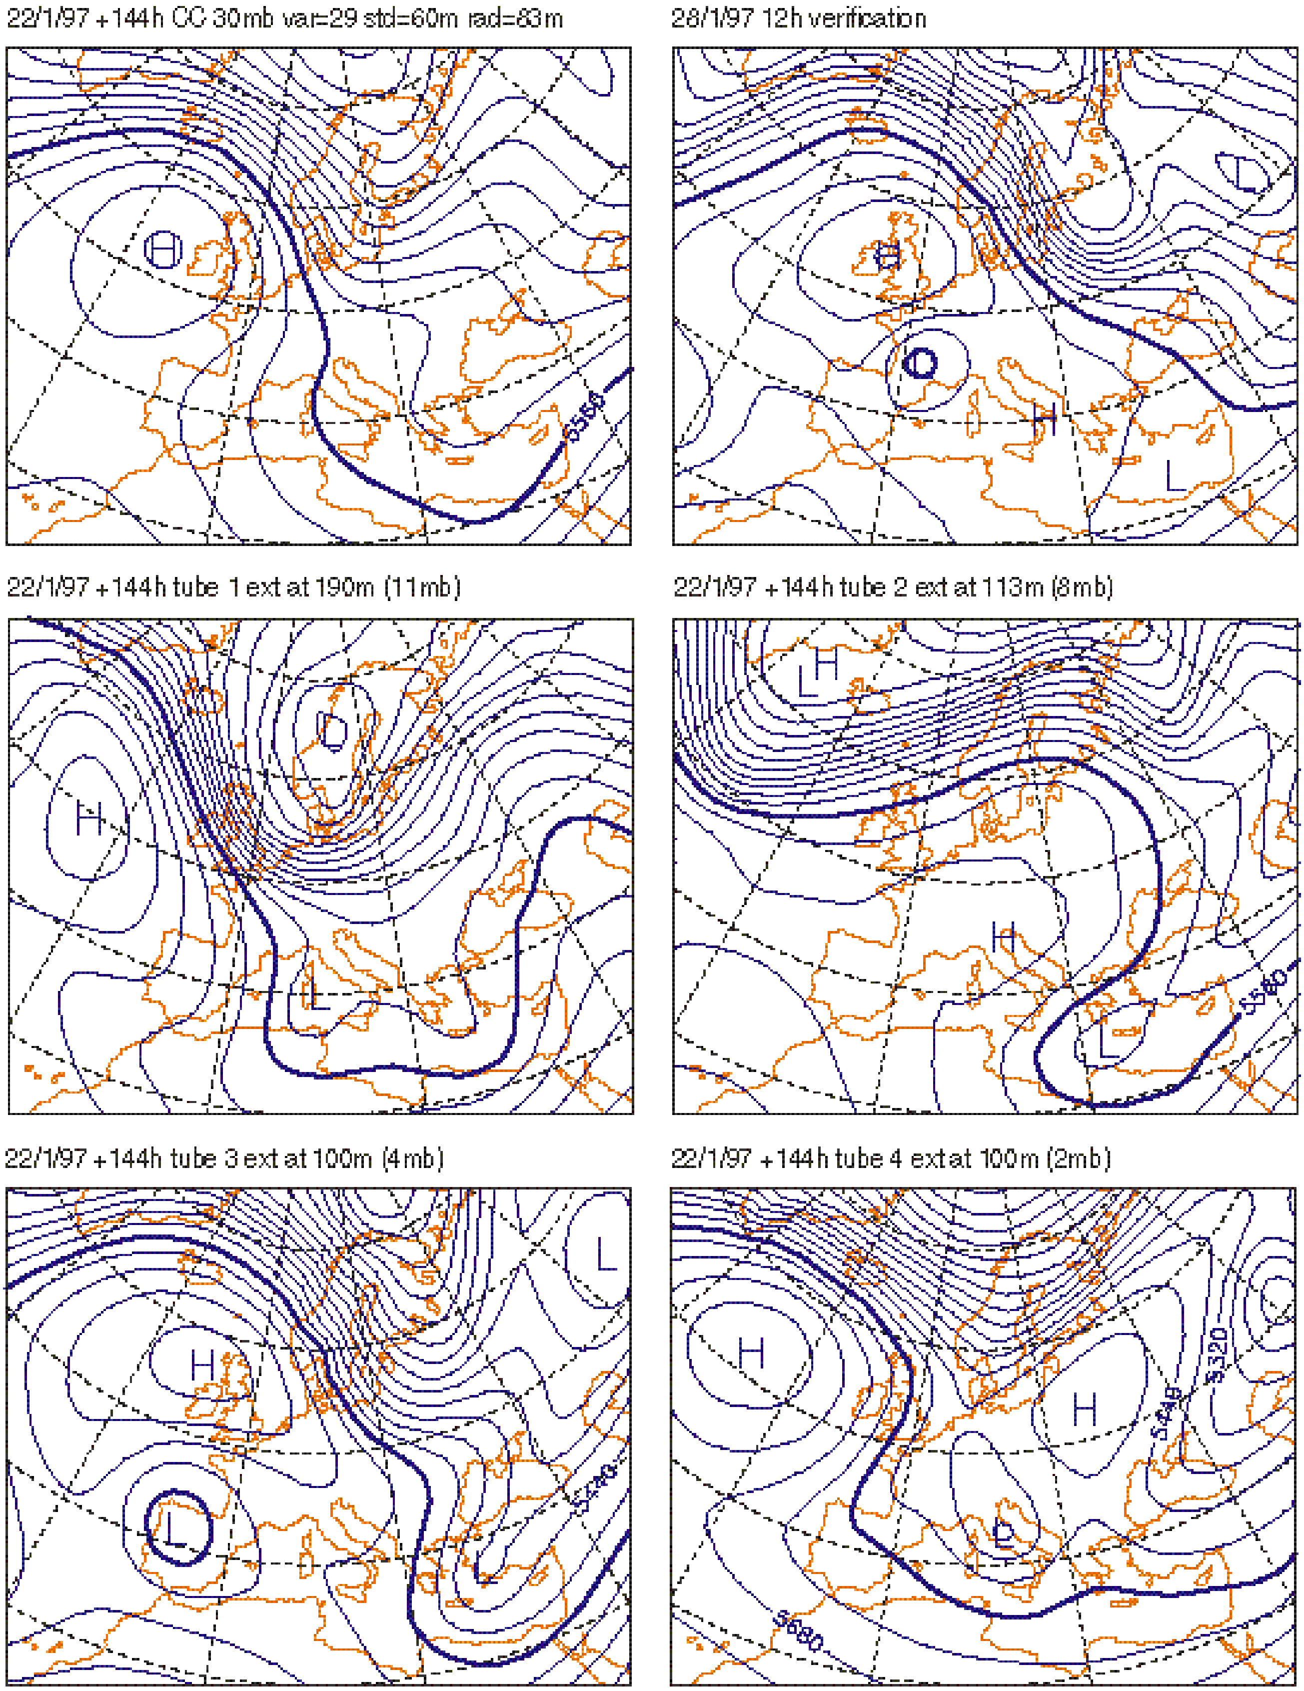

***Figure 5:** Tubing applied on the EPS based on 22nd January 1997, +144h. The central cluster standard deviation was limited to 60 m. Western Europe, 500 hPa height.*

the first tube tendency. The tube 3 and tube 4 tendencies are not represented by the clustering classification.

The clusters look more similar than the tubes extremes. This similarity, commonly reported by forecasters, is not only due to the 'smoothing effect' when averaging different fields (although it should not have such an impact if clustering only reflected 'natural' modes) but also to the distribution being monomodal, leading the centroids to lie around the ensemble mean rather than along the edges of the distribution (see section 4).

Probabilities

Grid-point probabilities have been computed from the EPS distribution to illustrate how the tubing product can be interpreted by operational forecasters to complement probabilistic forecasts, by identifying the main possible meteorological variants. Significant rain probability (precipitation > 1 mm/24h) ranges approximately from 10% to 60% over France (figure 8). Except over the Pyrenees region, where Northerly flows are likely to cause orographic precipitation, the risk of rain over the Western half of the country does not exceed 30-40%. The main synoptic pattern indicated by the central cluster leads indeed to dry anticyclonic conditions over this area. The 30-40% probability is mostly associated with the tendency represented by the tube 1: if the ridge does not develop as much as indicated by the main option, a secondary wave is likely to pass through at some stage, leading to a rainy interval over the considered area.

Frost probabilities (T2m < -5°C some time from +120h to +168h) indicate a 20 – 30% risk over Northern Germany (figure 9), although a deterministic forecast based on the central cluster synoptic pattern would probably exclude it: wet, cloudy weather is more likely to occur over this area with the North-Westerly cyclonic flow. The risk of frost is obviously associated with the tendency pointed out by the tubes 2 and 4, indicating a stronger ridge developing further over the continent and inducing anticyclonic, dry conditions over this part of Europe.

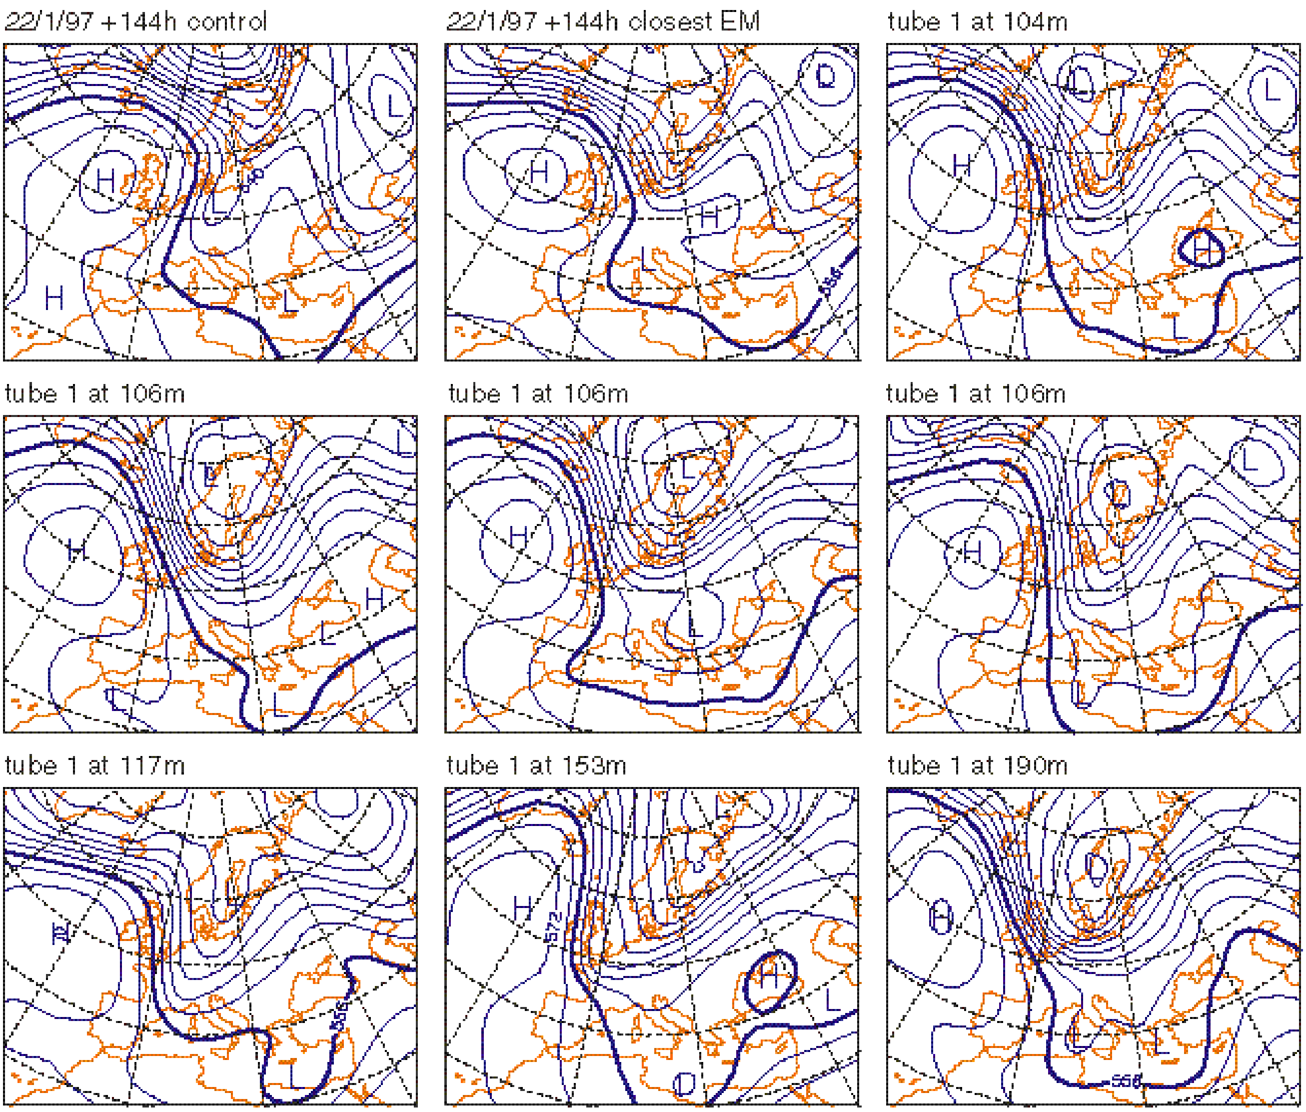

Figure 6: Ensemble forecasts grouped along the axis of the first tube.

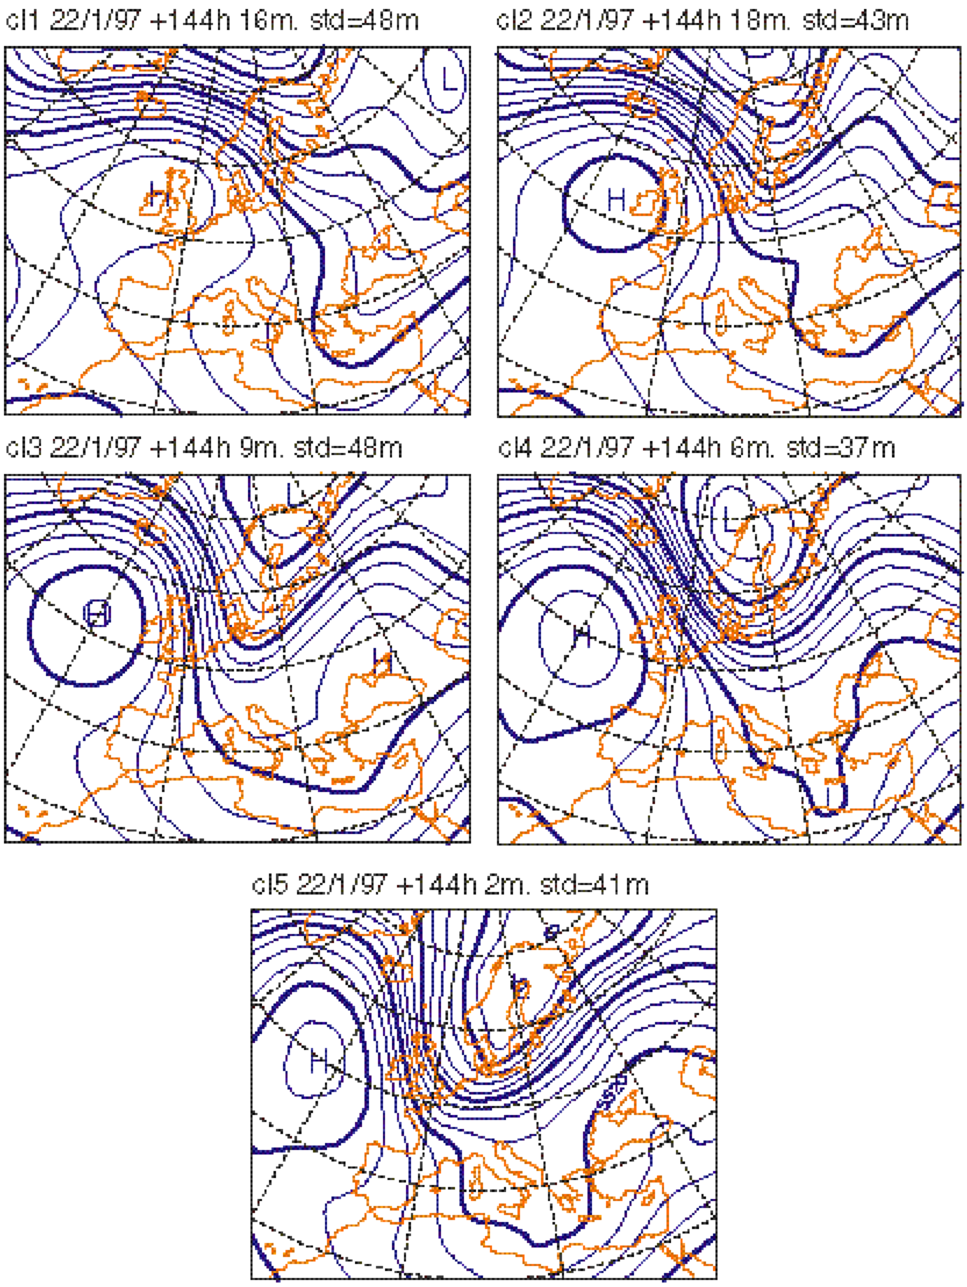

Figure 7: *Same as figure 5, but for a Ward clustering.*

Summary

The tubing method has been designed to facilitate a human interpretation of the distribution of meteorological forecasts produced by the EPS, in order to complement the probabilistic forecasts generated from weather parameters probability distributions. Tubing condenses the information from the EPS by highlighting the main differences between ensemble forecasts. The forecast distribution is represented by (i) the central cluster mean, grouping those ensemble forecasts lying around the ensemble mean, indicating the most likely meteorological pattern, (ii) the extreme forecasts located on the edges of the distribution, whose differences from the central cluster mean indicate the possible meteorological deviations from the most likely forecast.

Unlike conventional clustering methods, tubing is not designed to pick out alternative scenarios from a multimodal distribution. The underlying assumption is that ensemble distribution is generally monomodal, and distribution modes, when they exist, are spurious effects of sampling. The overwhelming performance of the ensemble mean, whatever the forecast distribution, is the basis on which the method was built. Still, tubing is an effective method of classification even when distributions are multi-modal, although it obviously emphasizes in this case the central part of the distribution at the expense of the existing modes, merely represented as possible deviations of the main central cluster pattern.

Tubing is indeed a method for *visualising* the forecast distribution, rather than a way of classifying arbitrarily ensemble members. One could object that the construction of the tubes assumes somehow an ‘axial’ multi-modality around the ensemble mean. As a matter of fact the tubes only give a condensed representation of the outer,

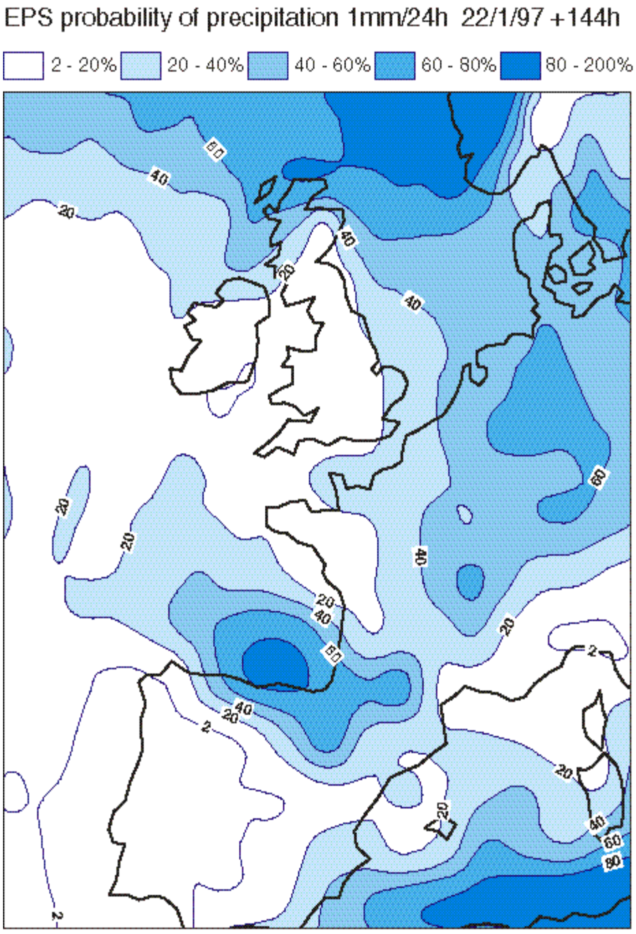

***Figure 8:** EPS probability of precipitation above 1 mm/24h, 22nd January 1997, +144h.*

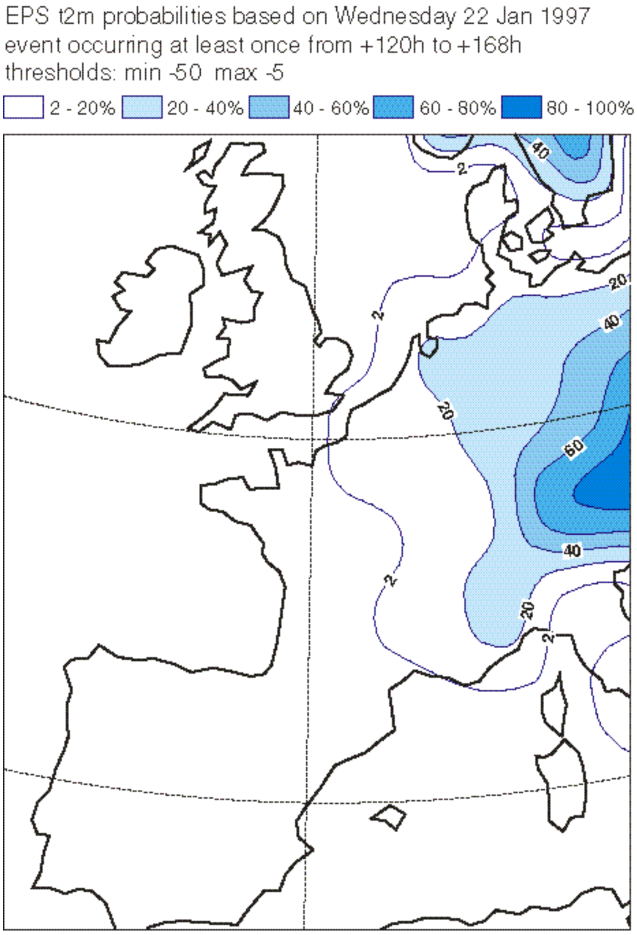

***Figure 9:** EPS probability of 2 m temperature below -5°C, 22nd January 1997, +120h to +168h.*

sparse part of the distribution: as soon as a direction is sampled by an ensemble forecast outside the central cluster, a tube will represent this direction or a close one. Nevertheless, the question of the significance of all these directions may be addressed. Empty, long tubes – i.e. containing only the extreme or, much less frequently, several members grouped together around the extreme – are notably questionable, because they artificially 'fill the gap' between the central, dense part of the distribution and one or several isolated members, somehow assuming the distribution to be continuous when there is no evidence nor theoretical reason for that. If the distribution modes were indeed reliable – which they are not on average – tubing should be avoided in those situations when the distribution is clearly discontinuous. ❑

Frédéric Atger

Summary of ECMWF Technical Memorandum 225

Simulation of fog with the ECMWF prognostic cloud scheme

João Teixeira

A prognostic cloud scheme has been operational at ECMWF since 1995. In this paper the performance of the cloud scheme in simulating fog is assessed. A case study is performed with the one-column version of the ECMWF model in order to analyse how the model reproduces the main mechanisms of fog generation and dissipation. Aclimatology of the model's fog is produced using data from the ECMWF Re-Analysis (ERA) and compared with climatological data. High resolution (T213L31) operational forecasts with the prognostic cloud scheme are compared with synoptic reports of visibility for Europe. In this context the relation between fog and temperature, on the one hand, and fog and wind speed, on the other, is explored in some detail.

The comparison between the simulated and the observed visibility, in the one-column model case study, shows that the evolution of fog is properly simulated. The analysis of the behaviour of the different parametrized physical processes suggests that the subtle balance between the various processes, fundamental for a realistic fog simulation, is achieved. This study also shows that in radiation fog a cooling process like radiation, that only changes T without changing q is rather essential in order to reach saturation. The comparison between the model's fog climatology and the climatological data shows that the model is able to reproduce most of the major fog areas, particularly over the ocean.

The analysis of the results from high resolution fog forecasts for Europe and the comparison with synoptic observations of visibility shows that the model simulates realistically some of the fog over Europe, particularly at 00 UTC and 06 UTC. At 12 UTC the fog in the model is usually underestimated. In general, the number of situations where the visibility is between 2 km and 11 km is underestimated by the model. It is shown that the model is always slightly closer to saturation than the observations and this might be helping the model in producing fog. The model reproduces the observed behaviour of fog as a function of temperature and wind speed rather realistically. ❑

Summary of ECMWF Technical Memorandum 234

On spurious chaotic behaviour in the discretized ECMWF physics scheme

Peter A E M Janssen and James D Doyle

Awell-known problem in atmospheric models is the large difference in time scales that may exist between dynamical processes and physics. Typically, physical processes such as turbulent diffusion near the surface may have time scales of a few minutes, while dynamical processes are integrated with a time step of 15 minutes. With the introduction of the semi-Lagrangian scheme the dynamical time step may become even larger. Therefore, in order to prevent numerical instability, physical processes are integrated in time with a semi-implicit scheme.

An example of an implicit scheme is the one of Crank and Nicholson, but owing to the nonlinearity of the physical processes only crude approximations to the Crank-Nicholson scheme have been implemented in numerical weather prediction models such as the ECMWF model. As a result there is no guarantee that the actual integration scheme for physics is numerically stable and this may have resulted in the occurrence of noise in quantities such as the heat flux, in particular since the introduction of the prognostic cloud scheme.

In this Technical Memorandum we analyse a simple one-dimensional evolution equation which reflects a relevant physical problem. Results from this example suggest that as long as a physical process has a short time scale and is sufficiently nonlinear, its discretized implementation will have properties akin to a well-known one-dimensional map studied by Feigenbaum. Thus, by increasing the integration time step sufficiently the discretized physics may show chaotic behaviour. However, the chaotic behaviour of the numerical scheme should be regarded as spurious because the continuous system is well-behaved. On the other hand, our analysis provides an explanation for the occurrence of noise in a number of physical quantities in the ECMWF physics.

It is clear that one should introduce a numerical scheme that avoids the above-mentioned spurious chaotic behaviour and a secure way to do this is to use a scheme that is always stable. A practical scheme that meets this requirement is a linearised version of the Crank-Nicholson scheme, and has been used in the wave prediction WAM model since 1985. However this scheme requires derivatives of physical processes which are not (yet) determined in the ECMWF model. We also discuss the idea of calling the physics subroutines (e.g. vertical diffusion) more frequently during one dynamical time step. This iterative approach may readily be implemented in the framework of the ECMWF atmospheric model and we have studied its impact on wind and pressure fields of the record low of January 10 1993. The iterative scheme gives a better representation of the low level jet in this case, while the consistency between surface wind and surface stress also improves.

We have also found that the iterative scheme for the vertical diffusion results in the elimination of spurious oscillations, which as a consequence may have impact on systematic errors in the atmospheric model. We thought therefore that it was of interest to perform a few 120 day runs in order to see whether there was evidence of any systematic differences between the control run and runs with the iterative scheme. Considerable differences in the mean 850 hPa temperature are seen over a number over areas such as North Africa. Similar changes in temperature were seen in a 120 day run in which the integration time step for dynamics and physics was reduced by a factor of three.

The iterative physics scheme for physics may therefore have a considerable impact on systematic errors of the atmospheric model. This iterative scheme was implemented operationally in December 1997. It should be realised, however, that this is most likely not the end of the story, because the iterative approach does not guarantee numerical stability. Therefore the introduction of a stable numerical scheme for physics, such as the linearised Crank-Nicholson scheme, seems most desirable. ❑

Summary of ECMWF Technical Memorandum 235

Use of the Fujitsu VPP700 for weather forecasting at ECMWF

D. Dent, G-R. Hoffmann, P.A.E.M. Janssen and A.J. Simmons

The introduction of the Fujitsu VPP700 into production use at ECMWF in September 1996 signalled a significant change from the previous dependence on shared memory parallel computers. The memorandum gives an overview of the forecasting activities, including the T213L31 semi-Lagrangian model, the variational data assimilation, the ensemble prediction system and the ocean wave forecasting.

Since the technical development necessary to be able to execute with good efficiency on a distributed memory platform was substantial, aspects of implementation and computational performance of the model on the VPP700 are discussed.

Operational running of a radically different platform posed new challenges in order to achieve reliable operational forecasting and good computer utilisation. Some early experiences are presented.

Finally, future developments in the many areas of research and operational forecasting are outlined. ❑

Summary of ECMWF Technical Memorandum 237

A modified set up of the advection scheme in the ECMWF wave model

Jean-Raymond Bidlot, Peter Janssen, Bjorn Hansen and Heinz Gunther

When the high resolution wave model was set up in late 1996, it was discovered that there were problems regarding excessive shadowing behind islands. This resulted in large zones for which the wave energy was largely underestimated. At the time, it was decided to implement the model operationally without the small islands that would otherwise be resolved with the new high resolution model (55 km). North south propagation of swell energy was also know to be limited to unrealistic narrow bands which are usually referred to as cigars.

These shortcomings of the wave propagation scheme could be addressed in many intricate manners, however a simple redefinition of the discretized propagation directions was sufficient to alleviate the previous problems. Specifically, the mean directions of the directional spectral bin were rotated by half the bin width. With this arrangement, the discretized propagation directions do not coincide with the two axis directions of the propagation scheme, and information can always be advected from all four neighbouring grid points. This feature is most needed as waves propagates around small islands or along the directions of the grid axis.

It was therefore possible to reintroduce the small islands that had been removed from the previous operational global wave model. Comparison with the old scheme shows a local improvement in the representation of the shadow zone behind islands as well as for northsouth propagation of swell. Global scores are fairly neutral, as would be expected from the inherent limitation of the propagation scheme. The benefit was nevertheless sufficient to justify the operational implementation of the new scheme. The first operational analysis produced with this new scheme was carried out for the 14th of May, 1997, 12 UTC. ❑

A Pinch of SALTT

Code optimisation on VPP700

The VPP700 is a ‘vector’ computer, which means quite simply that the high execution rates we expect can only be achieved by ‘vectorising’the code. Mostly this is an automatic process which the Fortran compiler is very skilled at performing. Users of ECMWF computers are well used to this process and have learnt techniques which make the vectorising process easier for the compiler. Success is usually measured by achieving a high vector CP percentage in the job accounting summary. Failure means that computations are performed in ‘scalar’ mode which is very much slower than the vector computational rate. Hence, the vectorising process is the single most important stage in optimising for good performance.

The next stage of optimising would be to look at ‘vector length’ information, since there can be enormous differences in performance between very short vectors(of order length 10) and longer ones (of order 100 or more). This information can be retrieved on a subroutine and even on a loop basis by using the Fujitsu profiling software (VPP_STATS=8 for example) and studying the trace files so generated using MPTOOLS. The available statistics include average, minimum and maximum vector lengths for each Fortran DO LOOP.

Since the execution of IFS (ECMWF's Integrated Forecast System) occupies such a large fraction of the VPP resources, it is interesting to take the investigation one stage further and try to find out more about the quality of the vector code being executed. Arecent study has investigated the use of a new tool for this purpose, a Simple Applications Level Tuning Tool (SALTT). This article discusses its use on IFS and comes to some preliminary conclusions.

An introduction to SALTT

The VPP700 hardware contains a set of counters and switches which can be labelled "The performance analyzer hardware". Using this, many different sets of detailed hardware events can be counted.

Recently, an easy to use software interface (SALTT) to the performance analyzer hardware has been provided. This is invoked by using the compiler 'hook' function, whereby a pair of calls is generated by the compiler at subroutine entry and exit time. In general, any code may be executed by these 'hook' calls; for example, they have been used in the past to trace the memory usage of IFS.

Using this interface, the contents of the appropriate hardware counters are extracted. Currently, only the hardware counters for the master PE of a parallel application can be interrogated.

Mechanism

There are two complementary ways to control what is measured in a run:

- by choosing which routines are compiled with the hook option (the main program MASTER must be included). Note that analysis of the results is simplest if only a small number of routines are included;
- by means of an environment variable (VPP_SALTT=n) to control extraction , where
  - n=0 gives an overall summary
  - n=4 extracts statistics for vector ADD, MULTIPLY and DIVIDE pipes
  - n=6 extracts vector LOAD statistics
  - n=7 extracts vector STORE statistics.

The measurements reported here have been obtained by including:

- MASTER on its own;
- MASTER + CNT4 (CNT4 is a high level controlling routine which excludes the IFS start-up costs);
- MASTER + one other routine, selected from the top 10 most expensive routines as measured in a profile run (see Table 1).

Subroutine	Category	%Total	%Vector	%Scalar	Vlen
cloudsc	physics	6.7	98.6	1.4	713
cuadjtq	physics	6	94.2	5.8	470
laitqm	SL	5.5	100	0	254
vdfexcu	physics	4.3	99.8	0.2	494
larche	SL	4	99.5	0.5	510
laitqmh	SL	4	99.7	0.3	253
cuascn	physics	3.9	98.4	1.6	587
lwttm	radiation	3.5	99.4	0.6	490
sgemmx	transforms	3.4			
lascaw	SL	2.7	81	19	730

Table 1: *T213 Profile (top 10 only)*

N.B. Vlen is the vector length as seen by the hardware and averaged over all vector operations within the routine. The vector register file is reconfigurable by the compiler within the range (128, 1024). In general, Vlen is less than the Fortran DO loop limit.

What is measured?

Each VPP700 PE contains a 'Vector Unit' (VU) and a 'Scalar Unit'. One of the hardware statistics measured is 'Vector Unit Busy'time. The vector unit is busy whenever any vector pipe is operating, so it is important to have a breakdown of the vector pipe activity. To achieve maximum performance rates, the ADD and MULTIPLY pipes should be simultaneously working (e.g. as in a matrix multiply). The tool reports how close the execution is to this ideal.

For the summary measurements (n=0) only the 'headline' performance is recorded, so it is easy to see the difference between the explicit performance defined by (+*/) in Fortran and the total vector unit computations.

For n=4, all floating point operations are recorded, including those involved in conditional arithmetic (the IF test itself) and all branches of the IF tests. Any address computations such as those necessary for indirect memory references will also cause the vector unit to be busy.

The memory references are subdivided into 'stride 1' (Vector memory block load or store) and 'non-stride 1' (Vector memory random load or store).

For the purpose of evaluating the processor performance, message passing activity only confuses the picture. Therefore nearly all of the IFS performance tests have been done using a resolution of T106L31 as this fits into a single (2 GB) PE memory. To confirm that the performance at T106 is reasonably close to T213, a small number of measurements have also been made using the production resolution (T213). By avoiding the setup phase of IFS, the execution rates are representative of any multiple of a 3 hour forecast period. Computational intensities have been computed (flops/memory references) on the basis of double precision real computations (one word = 8 bytes).

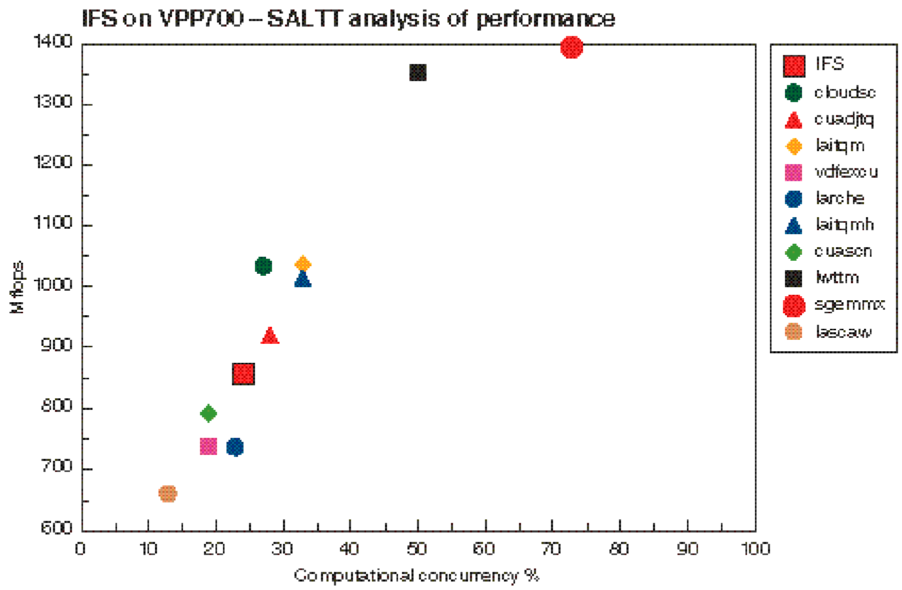

Figure 1: *Performance related to add/multiply overlap.*

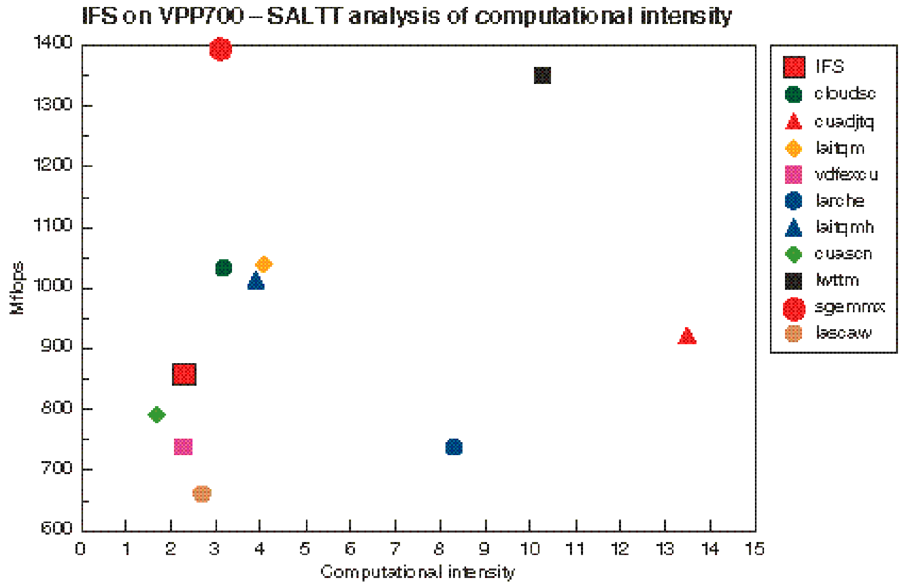

Figure 2: *Performance related to computational intensity.*

Figure 3 (below): *Performance related to add/multiply overlap and the way in which memory is accessed.*

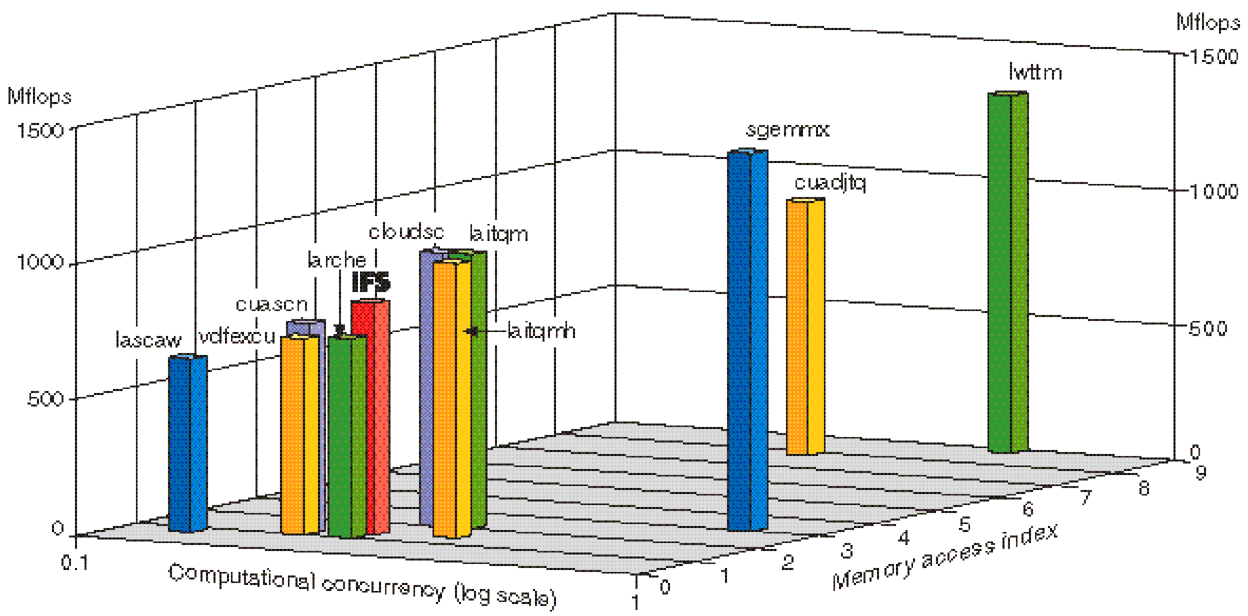

Summary of measurements

The measurements reported here relate to IFS configuration 1 at cycle 18R3, run in its operational form for 24 steps (6 forecast hours) but with all post-processing disabled.

T213 with start-up effects removed

800 Mflop/s
230 Mword/s loaded
126 Mword/s stored
Computational Intensity (CI) = 2.2

Total floating point operation count = 5 x 10^{12} for a one day forecast.

The sustained flop rate for a 16 PE execution is 12.8 Gflops. An estimate for the sustained rate on the full machine (105 PEs) is 63 Gflops (75% efficiency compared to 16 PEs).

A 16 PE run has about 94% parallel efficiency, so the degradation of floating point rate compared to the serial T106 execution is reasonable.

T106

857 Mflop/s
213 Mword/s loaded
139 Mword/s stored
CI=2.3

This is the ‘best’ execution rate seen since it is free from model start-up and message passing

Routine level statistics

When the SALTT statistics are studied for individual subroutines, the very large quantity of data available is difficult to assimilate. However, one can make general statements about some of the routines, for example:

(a) LWTTM has a very high execution rate (1350 Mflop/s) thanks mainly to extensive use of the fast square root function
(b) LASCAW does rather poorly (660 Mflop/s) due to substantial indirect addressing.

Since the peak speed of the VPP700 (and indeed of most high performance computers) depends on being able to perform ADD and MULTIPLY simultaneously, one of the statistics output by SALTT ‘Vector ADD / MULT pipes concurrently busy’ should be a good indication of the floating point performance of that routine. If this statistic is plotted for the ten most expensive IFS routines, a good relationship is seen with performance (Fig. 1). However, this takes no account of the way in which memory is being used. To study this aspect, we look at how much memory activity there is in these routines (computational intensity) and whether memory is accessed in a random fashion.

From Fig. 2 it can be seen that some routines with a high computational intensity perform well, while others perform poorly. The SALTT statistics give a measure of how well the memory load/store pipes are being used by reporting, for example: ‘Efficiency of load pipe as a percentage’. Ahigh efficiency will only be achieved when memory is accessed in a succession of neighbouring words (stride 1). Alow efficiency is usually an indication of significant random load/store (or non stride 1) access. To take this into account, we define a ‘memory access index’

MAI=(computational intensity) * (load pipe efficiency)

and plot this against computational concurrency and performance (Fig. 3). Code with a high memory access index either references memory sparingly or efficiently (or both). Now at least we can see that routines with a high MAI certainly do better than average (SGEMMX, CUADJTQ, LWTTM). Routines with a low MAI generally perform relatively poorly (LASCAW, VDFEXCU, LARCHE).

Conclusions

This performance measurement tool has proved to be very useful. Because the reported operations include those carried out inside multiple branch IF conditions, the headline flop rates may not be directly comparable with other platforms. However, the sustained rate when compared to the theoretical machine speed is very good.

The variations within key subroutines are substantial. Compared to previously reported values (ECMWF Newsletter No. 75 - Spring 1997), the rates are higher because of these over-estimates. Some of the relatively poor performances are easily explained by (for example) the extensive indirect addressing. In other cases, the non unity memory striding creates a performance penalty.

It is not obvious if further improvements can be made in IFS performance. It may be that some small application changes can be made, particularly in the choice of compiler technique for vectorising loops containing IF structures. It is also possible that the compiler generated code can be improved. Such possibilities can only be determined by much more detailed studies. However, the incentive to achieve single processor performance improvements remains high given the large number of IFS executions every day.

The computational intensity for the model as a whole is 2.2 which is much higher than that measured several years ago on a CRAY C90 (approximately 1). This is not due to the over-estimate of floating point operations in the conditional code since the same CI is measured when the model is run with physics disabled. Rather, it is likely to be due to the relatively numerous vector registers available on VPP. The CRAY compiler was forced to generate temporary vector store and load operations for many loops where the capacity of the C90 vector registers was insufficient. ❑

David Dent

Closure of ECFILE service

ECFILE is the user file archive service that has been in use since 1984. It runs on the IBM ES-9000 system, using CFS (Common File System) as its underlying storage manager. It has been superceded by the ECFS system (see ECMWF Newsletter 78, p. 8).

I would like to remind you that the ECFILE service will terminate on 31 July 1998. Therefore, users who wish to keep any of their ECFILE files must move them to the new service before that date. Once the ECFILE service has terminated there will be no possibility to retrieve ECFILE files.

Please note the following points:

- no automatic migration of user files from ECFILE to ECFS will be done, users must migrate their files themselves;
- a utility exists to help move individual files from ECFILE to ECFS, see the section "To copy files from ECFILE to ECFS" in the ECFS Users Guide;
- do not leave it to the last minute to migrate the files. If many users do this, the system will not be able to cope with the resulting copy load;
- anyone who has to migrate a large number of files (>500) from ECFILE to ECFS should contact User Support for further assistance;
- try to merge small files into one larger file on ECFS (e.g. using 'tar'or 'cpio'), as ECFS can handle one large file much more efficiently than many small files.

Finally, any queries about the change of data handling system should be addressed to User Support in the first instance. ❑

Andrew Lea

Announcement and Preliminary Call for Papers

Towards Teracomputing

Eighth ECMWF Workshop on the Use of Parallel Processors in Meteorology

ECMWF, Reading, England, 16-20 November 1998

Every second year the European Centre for Medium-Range Weather Forecasts (ECMWF) hosts a workshop on the use of parallel processors in meteorology. This year we will hold our 8th workshop in this well established series from 16 to 20 November 1998. The emphasis of this workshop will be on 'TeraComputing', that is, achieving sustained teraflop performance in a production environment.

Our aim is to provide a venue where

- users from our Member States and around the world report on their experiences in the field of parallel processing during the last two years. Plans for the future and especially the demands for computing power and the supporting infrastructure can also be presented;
- vendors of parallel supercomputers are able to talk about their current and future products to managers and end users of meteorological computer centres;
- meteorological scientists present their achievements in the development of parallel techniques and algorithms, and exchange ideas on the use of parallel supercomputers in future research;
- computer scientists give an update on their efforts in providing tools which will help users to exploit parallelism on supercomputers;
- the prospects and challenges of creating a computer centre infrastructure for 'TeraComputing' can be discussed. This topic has been included following a suggestion from the last workshop.

The workshop will consist of a limited number of presentations from invited speakers, plus a series of 20 minutes contributions. As before, the morning of the final day will be reserved for an open discussion session. The workshop will start on Monday 16th November at 9:30 am and close on Friday at 12:00 noon. It is planned to publish the proceedings of the workshop.

If you wish to contribute a talk, please submit a title and a short abstract of about 10 lines to ECMWF by 15 August 1998 .

The registration fee for the workshop will be 75.00 pounds sterling including refreshments, conference dinner and a copy of the proceedings. The registration fee will be waived for speakers.

Attendance at the workshop will be limited to around 100 persons. If you want to be included in the mailing list for further information, please contact by post, fax or e-mail:

Norbert Kreitz
European Centre for Medium-Range Weather Forecasts
Shinfield Park,
Reading, Berkshire
RG2 9AX
United Kingdom
Fax: +44 118 986 9450
E-Mail: NKreitz@ecmwf.int

ECMWF Calendar 1998

May 6 - 7	Policy Advisory Committee	*9th*
May 11 - 12	Security Representatives meeting	
May 18 - 19	Finance Committee	*59th*
Jun 29 - 2 Jul	Workshop - *Modelling and data assimilation for land-surface processes*	
Jun 15 - 16	Expert meeting on EPS	
Jun 17 - 18	Seasonal Forecasting Users meeting	
Jun 23 - 24	Council	*48th*
Sep 7 - 11	Seminar - *Recent developments in numerical methods for atmospheric modelling*	
Sep 28 -30	Scientific Advisory Committee	*27th*
Oct 7 - 9*	Technical Advisory Committee	*26th*
Oct 13 - 14*	Finance Committee	*60th*
Nov 2 - 4	Workshop - *Diagnosis of Data Assimilation Systems*	
Nov 9 - 13	Workshop - *WGNE/GCSS/GMPP - Cloud processes in large-scale models*	
Nov 16 - 20	8th Workshop on *The use of Parallel Processors in Meteorology - Towards Teracomputing*	
Dec 2-3	Council	*49th*

* Subject to change

Table of Member State TAC Representatives, Computing Representatives and Meteorological Contact Points

Member State	TAC Representative	Comp. Representative	Met.Contact Point
Belgium	Dr. W. Struylaert	Mrs.L.Frappez	Dr. J.Nemeghaire
Denmark	Dr. P. Aakjær	Mr. N.Olsen	Mr. G.R.Larsen
Germany	Mr. G-R.Hoffmann	Ms.E.Krenzien	Dr. D. Meyer
Spain	Mr. T. Garcia-Meras	Mr. E. Monreal Franco	Mr. F. Jiminez
France	Mr. D. André	Mrs.M.Pithon	Mr. D. André
Greece	Mrs. M.Refene Dr. G.Sakellarides	Dr. G.Sakellarides	Mrs.M.Refene Dr. N.Prezerako
Ireland	Mr. J. Logue	Mr. L.Campbell	Mr. M.R. Walsh
Italy	Dr. S.Pasquini	Dr. G. Tarantino	Dr. G.Maresca
Yugoslavia*			
Netherlands	Mr. S. Kruizinga	Dr. J.W. de Vries	Mr. G.Haytink
Norway	Mr. K.Björheim	Mrs.R.Rudsar	Mr. P. Evensen
Austria	Dr. G.Wihl	Dr. G.Wihl	Dr. H.Gmoser
Portugal	Mrs.I. Barros Ferreira	Mr. C.M. Fernandes	Mrs. F. Prates
Switzerland	Mr. H.Müller	Mrs. C. Raeber	Mr. R.Mühlebach
Finland	Mrs. K.Soini	Mr. T. Hopeakoski	Mr. P. Nurmi
Sweden	Mr. L. Moen	Mr. R.Svensson	Mr. O. Åkesson
Turkey	Asc. Prof.S.Sari	Asc. Prof.S. Sari	Asc. Prof.S. Sari
United Kingdom	Dr. A.Dickinson	Dr. A.Dickinson	Dr. M.J.P. Cullen

* At its 37th Session (December 1992) the Council decided that the telecommunications link between ECMWF and Belgrade would be terminated with immediate effect, and that henceforth no ECMWF documents would be sent to the Federal Republic of Yugoslavia (Serbia and Montenegro).