In this issue

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, Berkshire RG2 9AX, UK

Fax: . . . . +44 118 986 9450

Telephone: National 0118 949 9000

International . . . . +44 118 949 9000

Public Web site . . . . http://www.ecmwf.int

Member States' Web site. . . . . . https://wms.ecmwf.int

Front Cover

Forest damage in France following the storms of December 1999. Can the Ensemble Prediction System provide risk assessments of such severe weather events (see article on page 2)?

© Photograph – Joachim Bertrand, Securité Civile.

Editorial

The article on page 2 provides a forward look into future developments of the Ensemble Prediction System. The authors indicate that an important application is likely to be risk assessment of extreme weather events, examples cited are the storms that caused so much damage in France over the 1999 Christmas period, the tropical cyclone that produced devastating flooding in Mozambique earlier this year, and the severe snow storm over the east coast of the USA in January 2000.

The revised land-surface analysis scheme for the operational model is described on page 8. The new scheme produces improved analyses of the two-metre temperatures and relative humidities, and better soil moisture analyses using an optimum interpolation scheme.

The article on page 13 describes the concepts of the parallel-programming standard OpenMP, and illustrates experimental applications of the technique to the IFS forecast model code.

Changes to the Operational Forecasting System

An upgraded version of the model (Cycle CY22r3, atmospheric model identification number 196 in the GRIB headers) was introduced operationally on 27 June 2000. This version includes important changes of the parametrization schemes for the land surface, lying snow and sea-ice.

Other changes included in this model version are:

- New (RRTM) long-wave radiation scheme;
- Improved ozone model;
- Improved treatment of precipitation processes in the first time-step;
- Use of more TOVS/ATOVS data (HIRS-12, AMSU-14; less constraint on AMSU-8; more off-nadir data);
- Use of actual buoy heights;
- Revised snow analysis;
- Revised observation and background error variances in 4D-Var;
- Use of digital filter for the gravity-wave constraint J_c in 4D-Var.

In addition, enhancements have been made to the post-processing partly related to the new surface scheme and partly developed for application within the ERA-40 project. A major technical change has been made to the 4D-Var data handling through the implementation of purpose-built 'ODB' database software to deal with observations, replacing the old CMAFOC file format.

CY22r3 gives an overall improvement of upper-air scores. The cold bias in the screen temperature over spring snow conditions, in particular over Fenno-Scandia, has been removed but a moderate warm bias has been noticed in places.

François Lalaurette

The future of ensemble prediction

Ensemble prediction is a technique by which a forecast model is run several times from initial conditions which differ by amounts consistent with uncertainties in the initial state. Since its inception in 1992 (*Palmer et al. 1993)*, the ECMWF ensemble prediction system (EPS) has become an established part of operational forecasting at ECMWF. The original basis for the development of the EPS was the notion that single deterministic forecasts are not consistent with the "scientific method", in the sense that the result from any scientific prediction is not complete without an estimate of the likely error associated with measurement and other experimental inaccuracy (*Tennekes 1991*). It would be easy to estimate the likely error associated with a numerical weather forecast using some average error based on a large number of past cases. However, the atmosphere is chaotic, and that implies not only that forecast accuracy can be sensitive to small uncertainties in starting conditions, but also that the amplification of the initial uncertainty itself depends on the initial state (*Palmer 2000*).

The current configuration of the EPS is briefly described below, and examples are shown of how the EPS can quantify the risk of severe weather in the medium range, in circumstances where single deterministic forecasts fail. This ability to quantify the risk of severe weather implies that the potential economic value of the EPS is much higher than can be obtained from single deterministic forecasts.

The EPS has now entered a more mature phase of development and, as discussed below, there is now an important practical objective to guide its development - as a tool for quantitative risk management in weather-sensitive commercial/humanitarian applications. This maturity is beginning to lead to the consideration of more direct linkage between the EPS and specific user application models. This will make the value of proposed developments to the EPS easier to quantify.

A brief description of the EPS

At the time of writing, the ECMWF EPS comprises 50+1 integrations of the operational ECMWF model at T_L159L40 (120 km) resolution (*Buizza et al. 1998*). The initial perturbations are based on the dominant singular vectors (finite-time instabilities) of the forward tangent propagator between day 0 and day 2 (*Buizza and Palmer 1995*). The singular vectors

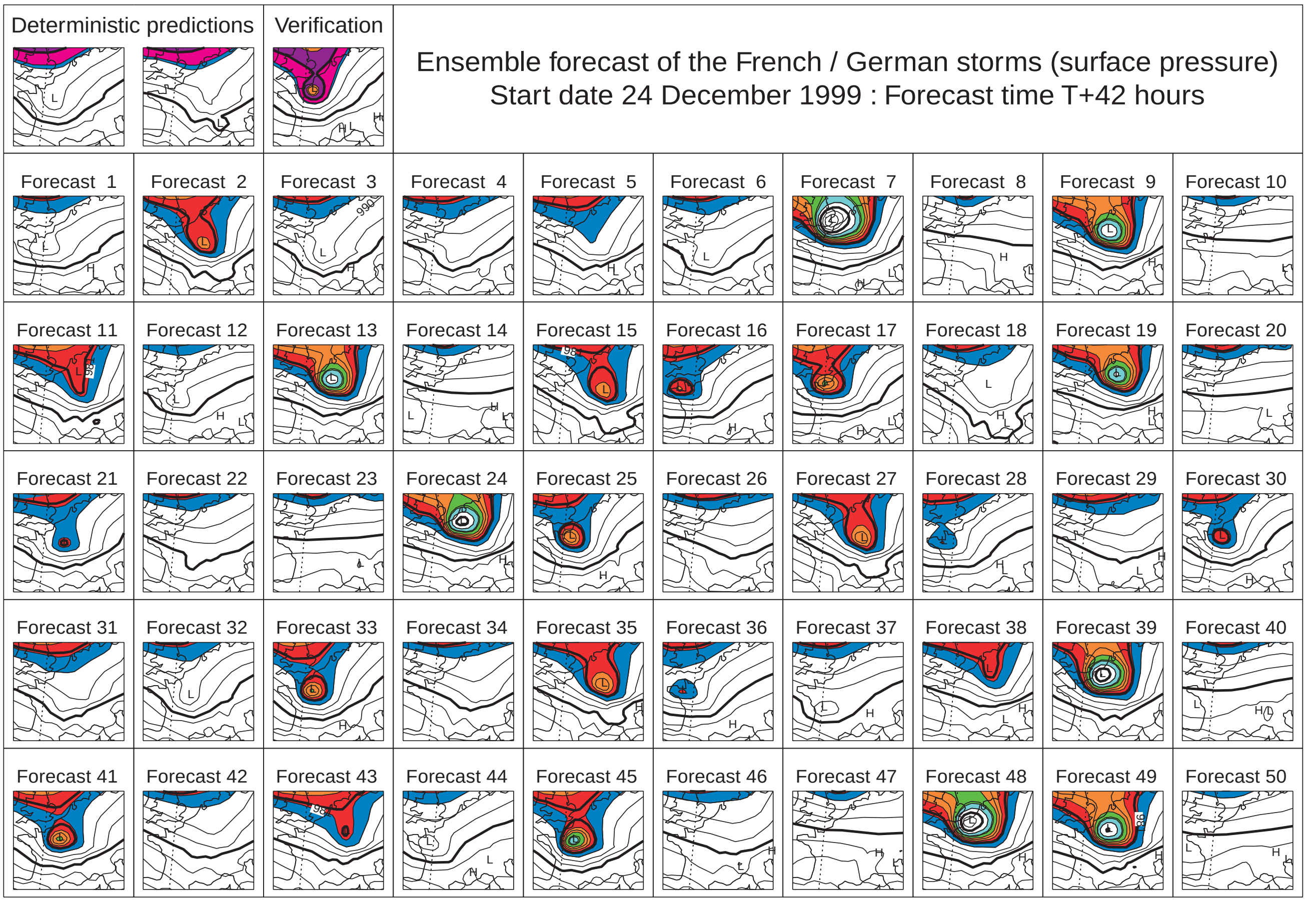

Figure 1 Stamp maps of surface pressure (coloured below 980 hPa) for a 42 hour forecast of the first Christmas storm at 06 UTC 26 December 1999. The T_L159 control forecast, the T_L319 operational forecast and the individual ensemble members are shown. Even though the single deterministic forecasts do not forecast the storm, the EPS shows that there is a significant risk of such a storm (indeed a very large risk by comparison with its climatological frequency).

are linearly combined taking into account the expected amplitude of analysis error (*Molteni et al.* 1996). The principal reason for using singular vectors is to "side-step" two problems. Firstly there are many unquantified assumptions in data assimilation and the true analysis-error probability density function (PDF) is, in practice, not well known; this means that a true random sampling of the initial PDF is impossible. Secondly, the number of samples possible is very much smaller than the number of possible choices of initial perturbation. Inadequate sampling will generally lead to an overly confident estimate of forecast reliability.

In addition to initial perturbations, the model equations in the EPS are perturbed using the stochastic physics scheme (*Buizza et al. 1999*). In this scheme, the physical tendencies are perturbed stochastically at each time step and grid point, representing random uncertainty in the formulation of the model equations.

It is proposed to increase the horizontal resolution of the operational EPS to T_L255 (80 km) in the second half of 2000. In addition to operational EPS results, some of the results below are shown for this higher resolution.

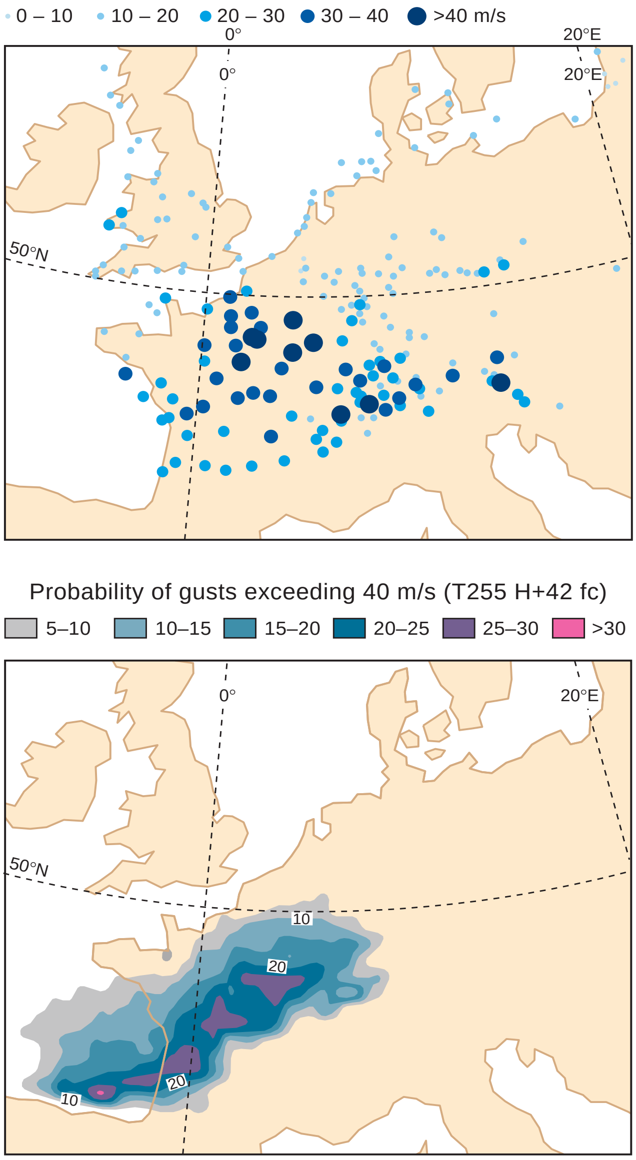

Figure 2 Observed gusts (top panel) and the 42-hour EPS forecast probability of gusts exceeding 40 m/s (computed from the 850 hPa wind speed, bottom panel).

Ensemble predictions of severe weather.

There is no doubt that National Meteorological Services are judged, by the public at large, by their ability to provide timely warnings of severe weather. By their nature, severe weather events are often associated with intense atmospheric developments on rather small scales. Such developments usually arise because of strong instability of the flow, suggesting that these forecasts, more than others, will be sensitive to small uncertainties in starting conditions.

In such situations, the guidance from single deterministic forecasts can be unreliable, and this is often manifest in day-to-day inconsistency of such forecasts. In this section, three case studies are shown demonstrating the value of risk assessment derived from corresponding ensemble predictions.

Figure 1 shows individual members of a 42-hour ensemble forecast from a T_L255 ensemble integration for the 26 December 1999 storm over France (in terms of surface pressure). Both the T_L319 operational forecast and the T_L159 control forecast failed to predict the storm at this range (the deterministic forecasts were more successful at both shorter and longer ranges, illustrating the point made above concerning the unreliable nature of single deterministic forecasts of severe weather). It can be seen that many members of the ensemble successfully simulated the storm. As such, the risk of such a storm was forecast with a probability vastly exceeding its climatological probability.

As an example of a possible forecast product, based on such an ensemble, Figure 2 shows the probability of gusts exceeding 40 m/s, based on a simple gust parametrization (that of mixing down air from 850 hPa; Ernst Klinker 2000, personal communication). This parametrization, applied to the ensemble predictions, forecasts a probability up to 30% of such damaging gusts over a swath of northern France. Six hours later (not shown) the region of maximum probability of such gustiness had moved on to Germany. This ensemble forecast, continued to 96 hours, also gave a significant probability of a second storm over France (not shown). The operational EPS (at T_L159 resolution, also not shown) also gave a warning of the storm. However, the gustiness product was somewhat weaker and less well positioned at this lower resolution.

A second example is shown in Figures 3 and 4. The severe snow storm over the US East Coast on 25 and 26 January 2000 was missed by the main operational models. The failure to issue a warning of the risk of severe weather was criticised in the media. The storm caused severe damage with loss of lives, mainly because of the intense snowfall associated with its passage.

Figure 3 shows some of the individual members of the EPS (in terms of 1000 hPa geopotential height and precipitation). It can be seen that the deterministic forecasts have the precipitation bands out to sea, whilst individual EPS members correctly develop an intense low-pressure system over land, with associated precipitation. It can be noted that the ensemble-mean forecast for this event is not

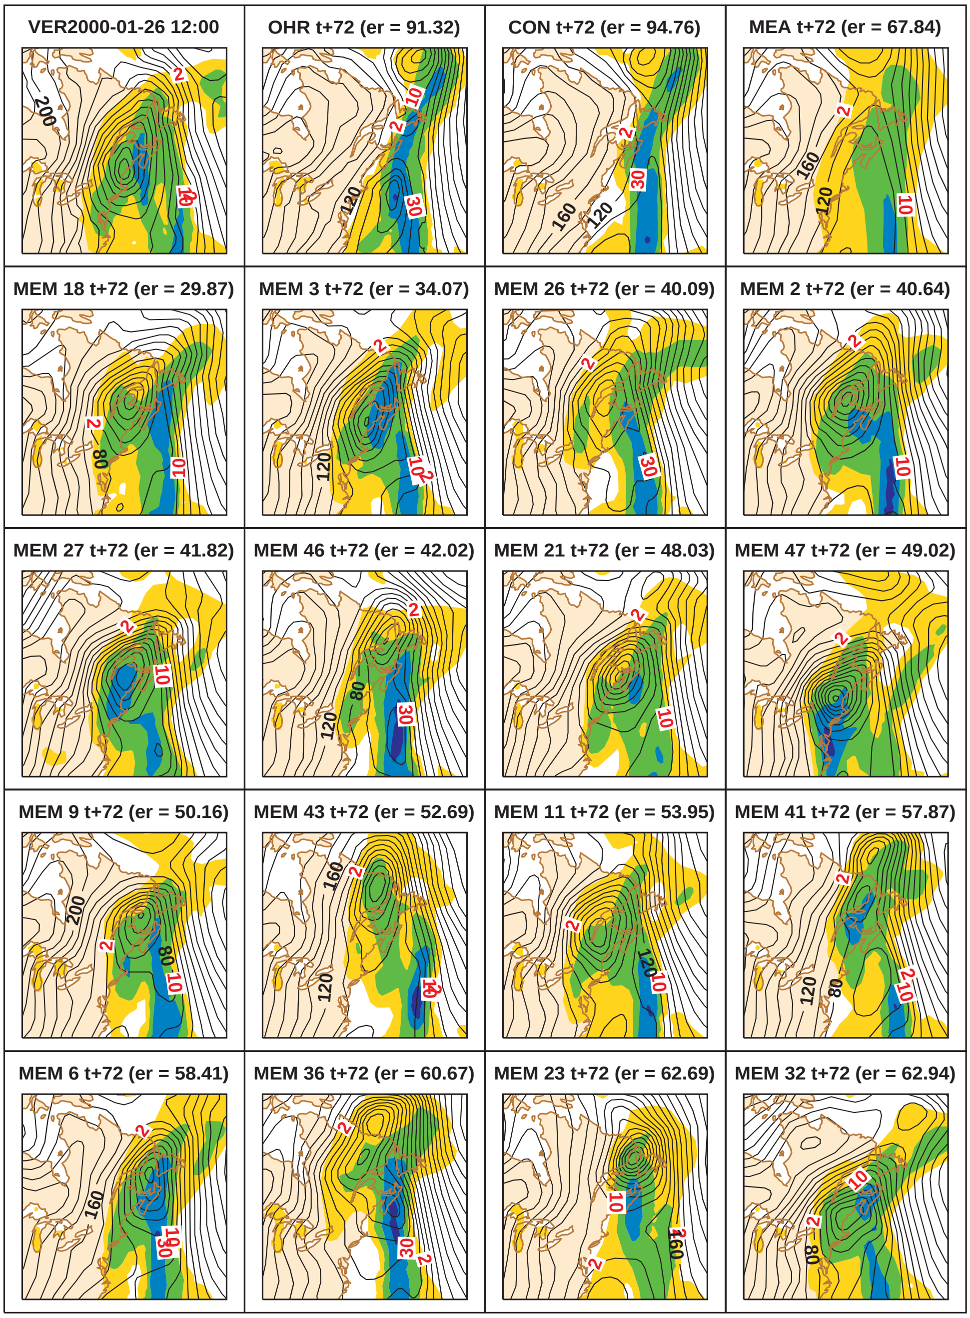

Figure 3 Postage stamp maps for 1000 hPa geopotential height and precipitation for the verifying analysis (first top panel), and for 72-hour forecasts started at 12 GMT on 23 January 2000 of the high-resolution model (second top panel), the EPS control (third top panel), the ensemble-mean (fourth top panel), and the first 16 EPS perturbed-members (subsequent panels) with the smallest root-mean-square-error for geopotential height inside a region centred around the observed cyclone. Contour interval is 20 m for geopotential, and contour isolines are 2, 10, 30 and 60 mm for precipitation.

itself successful. This should be taken as a warning that trying to use the ensemble to provide a modified deterministic (i.e. non-probabilistic) prediction is not itself a reliable procedure.

Figure 4 shows the 0-24 ECMWF operational forecast from 12 UTC on 25 January, which can be considered as a good approximation to the observed precipitation field (1 mm of water is equivalent to 1 cm of snowfall). The 72-hour forecast from the ECMWF operational model failed to predict intense snowfall over the land. By contrast, the 72-hour EPS forecast a probability of 10-to-60% (depending on the geographical location along the US East Coast) of more than 10 mm/d of precipitation, and 2-to-30% probability of more than 20 mm/d. The EPS also indicated a significant probability of enhanced 10 m wind speeds compared with the deterministic forecasts (not shown).

Figure 5 shows an example of an ensemble prediction of the position of tropical cyclone Eline which devastated Mozambique in February 2000 up to 5 days ahead from 18 February 2000. In this example, initial perturbations were made using tropical diabatic singular vectors (*Barkmeijer et al. 2000, Puri et al. 2000*). Such perturbations are not yet computed operationally, but it is hoped to introduce them into the EPS in 2001. The ensemble shows that there is a high probability that the cyclone will strike Mozambique, although at a range of 5 days, the precise position of landfall is uncertain (interestingly the ensemble shows a somewhat bimodal probability distribution of landfall, with the verifying analysis and the high resolution operational forecasts taking different modes).

The potential economic value of probability forecasts

It has been argued that single deterministic forecasts of extreme weather are, by their nature, likely to be unreliable. Ensemble forecasts, on the other hand, should be capable of estimating the risk of extreme weather more reliably. How can this be quantified? Consider a simple decision model used to estimate the potential economic value of weather forecasts (*Murphy 1977, Richardson 1998, 2000*). A user can suffer a loss L if a meteorological event E occurs and no precautionary action is taken. The loss is avoided if precautionary action at cost C is taken. The weather forecasts are used to decide when to take precautionary action.

Figure 6 shows the value of the EPS as a function of the user ratio C/L for the event E: precipitation greater than 10 mm/day (of water equivalent). Also shown for comparison are the values of two deterministic forecasts – the T_L159 control and the ensemble mean. Zero value means that the information provided by the forecasts is of no more use than information associated with a knowledge of the climatological frequency of E. A value of unity would imply a perfect deterministic prediction system.

Both of the deterministic forecasts provide a simple but unreliable criterion for deciding when to take precautionary action: take action when E is forecast; do not take action when E is not forecast. For the EPS, a more sophisticated decision strategy, making use of the forecast probability distribution can be adopted. A user with small C/L should decide always to take precautionary action, except when the probability of E is sufficiently small. A user with C/L close to unity should only take precautionary action when the forecast probability of E is sufficiently high. In general the user should take precautionary action when the probability of E is greater than C/L.

The ensemble-mean forecast gives very poor value, worse than the control, whilst the full ensemble value remains high. The reason for this is straightforward; the ensemble mean is a smooth forecast, and will consistently underpredict the more extreme types of events. On the other hand, probability forecasts made using the ensemble probability distribution prove much more valuable in forecasting the risk of this precipitation event.

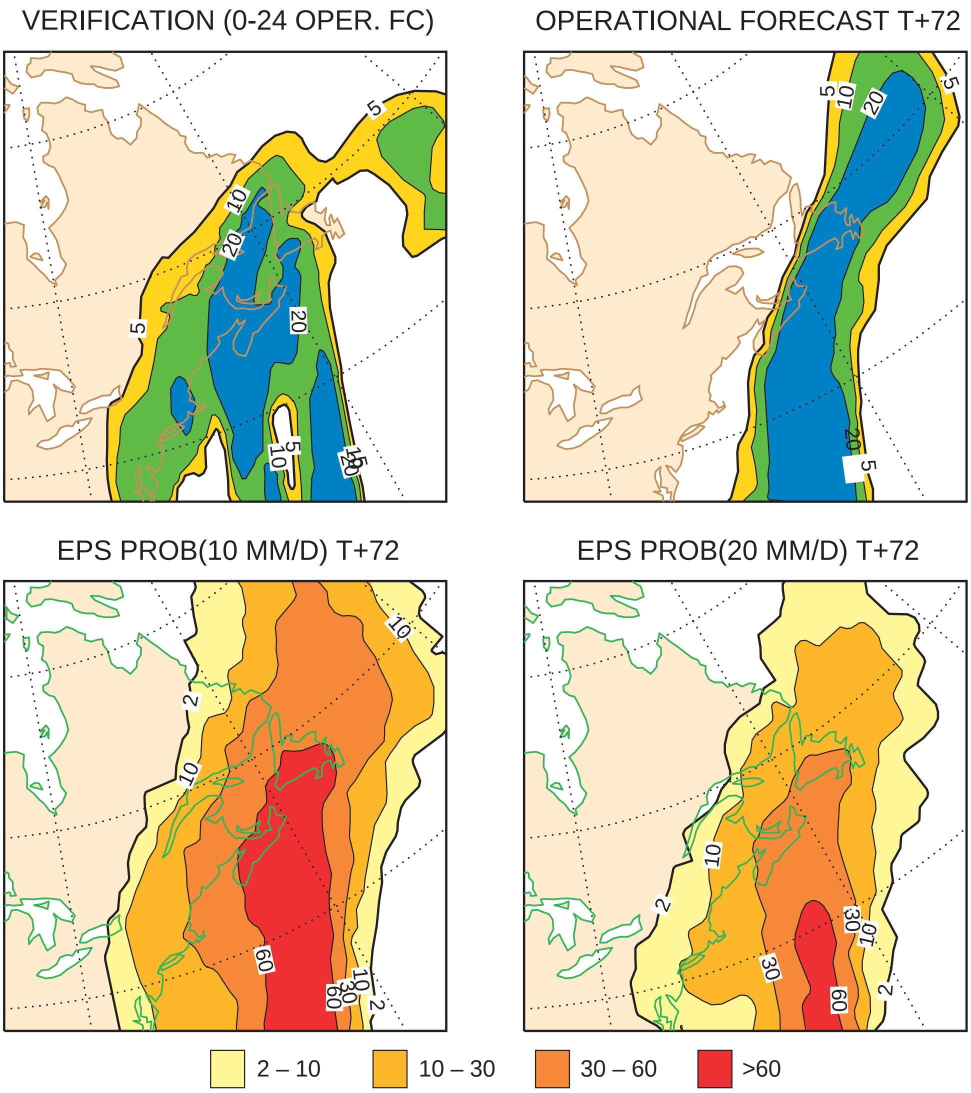

Figure 4 Verification field for precipitation (defined as the 0-24 hour high-resolution forecast started at 12 UTC on 25 January 2000, top left panel), 72-hour forecast from the high-resolution model (right top panel), and 72-hour EPS forecasts of the probability of more than 10 mm/d (bottom left panel) and 20 mm/d (bottom right panel) of precipitation. Contour isolines are 5, 10, and 20 for precipitation and 2%, 10%, 30% and 60% for probabilities.

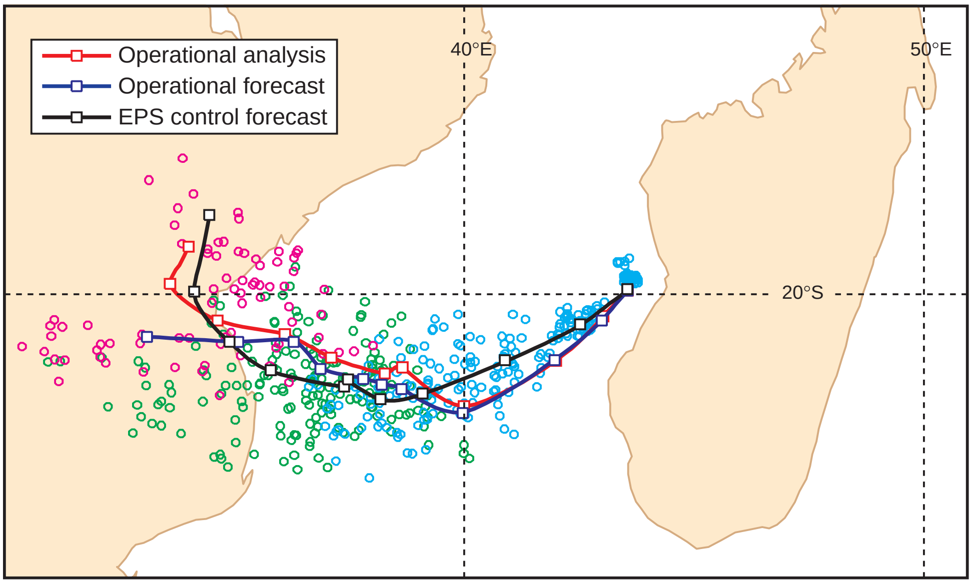

Figure 5 An experimental EPS showing the ensemble of tropical cyclone positions up to 5 days ahead for cyclone Eline which devastated Mozambique in February 2000. Symbols denote the cyclone position in the EPS 0-to-24 hour (blue), 48-to-72 hour (green) and 96-to-120 hour (purple) forecasts. The red line shows the observed path (square symbols identify the position every 12 hours), the blue-line the path predicted by the high-resolution model and the black line the path predicted by the EPS control.

EPS as a quantitative tool for risk management

The previous section has illustrated the value of the EPS as a quantitative tool for risk management; by contrast it has been shown that single deterministic forecasts are less valuable. In order to develop this notion of value, EPS output needs to be linked directly with specific user application models. Many examples of such user application models can be envisaged: prediction of damage from flood, storm, drought, prediction of electricity demand, prediction of the pay-out on some weather-related financial contract. A schematic of this notion is illustrated in Figure 7. Output from each member of the EPS is first passed through an empirical model to correct for model bias, and to apply down-scaling to give values appropriate to specific geographical points. For each member of the ensemble the corrected forecast data is fed into the user application model. The end product will be a probability distribution of damage/demand/pay-out. If the probability of abnormally high damage/demand/pay-out is sufficiently high, then the user can take appropriate action. The trigger for such action depends on the user, for very high potential loss, the threshold probability for such an abnormality might be rather low. In this way, the need for direct probabilistic forecasts of weather parameters is obviated, and questions such as "How do customers make use of probability forecasts?" are circumvented.

As an example of such a procedure, *Hoffschildt et al. (2000)* considered the problem of ship routing, computing over a season of forecasts, an ensemble of optimal ship routes from Brest to New York, based on the individual members of the EPS (see also *Janssen, 2000)*. In situations where the ensemble of ship routes is consistent, the master can plot his route with confidence. In cases where the ensemble of ship routes is broad there is uncertainty in the optimum route, and if he has this option, the master might be advised to stay in port, until a more clear-cut option is available. If this is not an option then, as shown by *Hoffschild et al. (2000)*, the most likely optimal route based on the ensemble can be chosen. Because of non-linearity in the prediction process, this "most likely" optimum route may differ from the optimum route provided by either of the deterministic forecasts. In general *Hoffschildt et al. (2000)* found that significant fuel saving could be had using the most likely optimal ship route defined by the EPS, compared with the single estimate of optimal ship route defined by the high-resolution T_L319 single deterministic forecast.

This type of "end-to-end" analysis of the EPS is still in its infancy. The true value of the EPS will not be realised until it becomes an established procedure. This may require a radical change in the way in which EPS data are used in Member State National Meteorological Services.

Computational Demands for the EPS

The EPS is computationally demanding; although, on the other hand, it is a perfect application for multi-processor supercomputers. There are four principal components that contribute to the cost of the EPS:

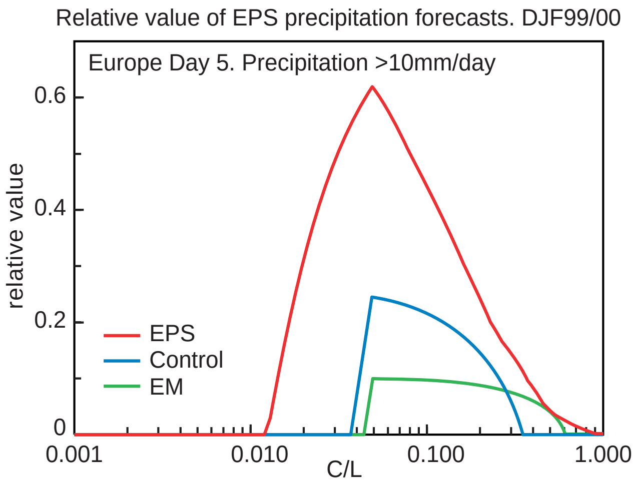

Figure 6 Potential economic value (see Richardson, 2000 for details) for forecasts of 10 mm/day, for a range of users with variable cost/loss ratios. Solid line: EPS, dashed line: control forecast, chain dashed line: ensemble mean.

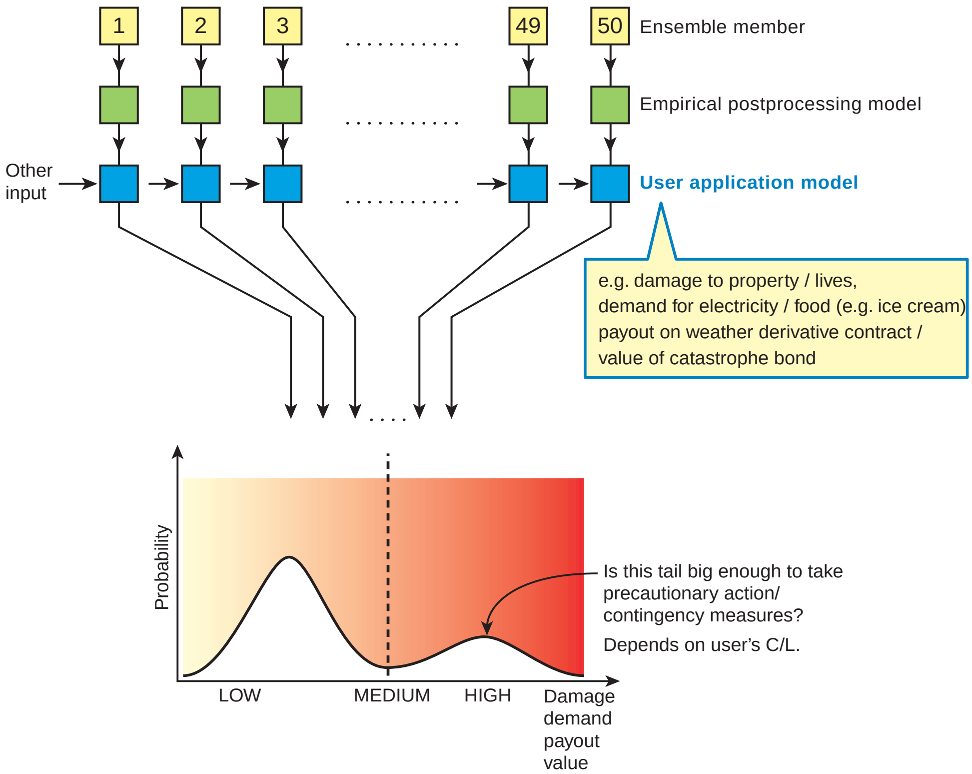

Figure 7 A schematic diagram showing that for quantitative risk assessment, the output of the EPS should be directly coupled to user application models. The output of such a system will be a probability forecast of user specific variables (damage, demand, pay out etc). From this, the user can directly assess if the risk of an abnormal situation warrants precautionary action.

- the construction of the initial perturbations;
- the resolution of the integrating model;
- the number of ensemble members;
- the length of the integration.

The construction of initial perturbations involves singular vectors, and this takes about 10% of the total cost of the EPS. The computation of diabatic singular vectors using the Hessian of the analysis cost function as initial metric increases the cost of this computation of singular vectors by about a factor of 4. Adding tropical singular vector computations would represent a further increase in cost. However, in view of planned increases in model resolution and ensemble size, there is no reason to suppose the computational cost for the initial perturbations will change significantly as a function of total cost, and this is therefore not the dominant concern when considering computer costs.

Balancing the need to have the best possible resolution for the nonlinear forecast model against providing a sufficiently large ensemble is a difficult judgement to make. The EPS is currently run at a resolution of about half of the high-resolution single deterministic forecast. However, as a tool for quantitative risk assessment of severe weather, it is clearly desirable to be able to run ensembles with a model whose resolution is such as to be able to simulate severe weather, otherwise the EPS will systematically underpredict the probability of severe weather. Insofar as T_L319 and T_L511 show a sensitivity of resolution to the simulation of severe weather, then these resolutions should be considered as eventual targets for the EPS. An intermediate increase in EPS resolution (to $T_L255L60$) is planned in late 2000, and documention of the improvement of the T_L255 ensemble over the T_L159 ensemble is currently in progress.

On the other hand, it is also necessary to ensure adequate ensemble size. For example, given the true forecast PDF, suppose that there is a 10% risk of some severe weather event (whose climatological probability may be orders of magnitude smaller than 0.1) – of the Christmas storm discussed above. Suppose that the EPS were to sample randomly from the true forecast PDF (which itself is problematic, see above), then with a 50-member ensemble one would expect 5 members to predict the event. Consider a particular 50-member ensemble in which no member predicts the event. Using a simple chi-squared test with a 1% confidence value it would be impossible to reject the hypothesis that such an ensemble could not have been drawn randomly from the true forecast PDF. This is an indication of inadequate ensemble size. By contrast, if the event was not predicted by any member of a 100-member ensemble, the chi-square test would indicate that this ensemble was not a random drawing of the true PDF.

One possible method for effectively increasing ensemble size is to run the ensemble more frequently, e.g. twice per day, instead of once. An optimal probability forecast could then be made by combining the two individual ensembles.

Finally, there is the question of integration length. In principle, providing the ensemble is able to produce reliable probabilities, then the EPS can be extended beyond 10 days (however, in practice, implementation of such a strategy should await a better simulation of the dominant intraseasonal modes of variability in the model - there is evidence to suggest that the representation of such modes may well improve with an interactive ocean). However, increasing the forecast range to 20 days will approximately double the cost of the EPS (it would more than double the cost if the model is run with an interactive ocean).

The EPS is undoubtedly computationally expensive, and the factors discussed above all increase its computational cost significantly. A crucial question is whether such an increase can be justified. This depends on what really is the value of the EPS to the ECMWF Member States. It has been argued that if the links between EPS output and application model input were more strongly developed, then the true value of the EPS as a quantitative tool for risk management would be realised. The development of this methodology will be a key to the development of the EPS itself.

References

Barkmeijer, J., **R. Buizza**, **T. N. Palmer** & **K. Puri**, 2000: Tropical singular vectors with linearised diabatic physics. *Q. J. R. Meteorol. Soc.*, to appear.

Buizza, R., & **T. N. Palmer**, 1995: The singular vector structure of the atmospheric global circulation. *J. Atmos. Sci.*, **52**, 1434-1456.

Buizza, R., **T. Petroliagis**, **T.N. Palmer**, **J. Barkmeijer**, **M.Hamrud**, **A. Hollingsworth**, **A. Simmons** and **N. Wedi**, 1998: The Impact of Model Resolution and Ensemble Size on the Performance of an Ensemble Prediction System. *Q. J. R. Meteorol. Soc.*, **124**, 1935-1960.

Buizza, R., **M.J. Miller** & **T.N. Palmer**, 1999: Stochastic simulation of model uncertainties in the ECMWF Ensemble Prediction System. *Q.J.R. Meteorol. Soc.*, **125**, 2887-2908.

Hoffschildt, M., **J.-R. Bidlot**, **B. Hansen** & **P.A.E.M. Jannssen**, 2000: Potential benefit of ensemble forecasts for ship routing. *ECMWF Technical Memorandum No. 287.*

Janssen, P.A.E.M., 2000: Potential benefits of ensemble prediction of waves. *ECMWF Newsletter 86*

Molteni, F., **R. Buizza**, **T.N. Palmer**, **T. Petroliagis** 1996: The ECMWF ensemble prediction system: methodology and validation. *Q. J. R. Met. Soc.*, **122**, 73-119.

Murphy, A.H., 1977. The value of climatological, categorical and probabilistic forecasts in the cost-loss ratio situation. *Mon. Weather Rev.*, **105**, 803-816.

Palmer, T.N., **F. Molteni**, **R. Mureau**, **R. Buizza**, **P. Chapelet** & **J. Tribbia**, 1993: Ensemble Prediction. 1992 *ECMWF Seminar Proceedings.*

Puri, K., **J. Barkmeijer** & **T.N. Palmer**, 2000: Ensemble prediction of tropical cyclones using targeted diabatic singular vectors. *Q. J. R. Meteorol. Soc.*, to appear.

Palmer, T.N., 2000: The prediction of uncertainty in weather and climate forecasting. *Rep. Prog. Phys.*, **63**, 71-116.

Richardson, D., 1998: Obtaining economic value from the EPS. *ECMWF Newsletter No. 80.*

Richardson, D., 2000: Skill and relative economic value of the ECMWF ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **126**, 649-667.

Tennekes, H., 1991: Karl Popper and the accountability of numerical forecasting. In New developments in Predictability. *ECMWF Workshop Proceedings*, ECMWF, Shinfield Park, Reading UK, 1991.

T.N. Palmer, J. Barkmeijer, R. Buizza, E. Klinker and D. Richardson

A revised land-surface analysis scheme in the Integrated Forecasting System

Two significant changes on the land-surface analysis were implemented in the operational ECMWF 4D-Var system during 1999. In March (Cy19r2), an analysis of two-metre temperature and relative humidity based on univariate optimum interpolation using SYNOP observations was put into the operational system. In July (Cy21r2), the previous operational soil moisture analysis based on a nudging scheme was replaced by an analysis of soil moisture and soil temperature based on optimum interpolation. These modifications have had a positive impact on the quality of ECMWF near-surface products and are also expected to improve the quality of products generated by the 40-year reanalysis project (ERA-40). The purpose of this paper is to describe and illustrate this revised land surface analysis scheme.

In August 1993, the relaxation of both the temperature and moisture content at the bottom of the soil layer to their climatological values was replaced by a zero flux condition for heat transfers and a free-drainage condition for water transfers. By summer 1994, the land surface had drifted to an excessively dry state over most of the Northern Hemisphere continents, with a detrimental impact on forecast scores. With no relaxation to climatology and no initialisation of the soil prognostic variables, nothing prevented the land-surface scheme from drifting to an unrealistic state. Such drifts are the consequence of a positive feedback with the atmosphere when the land-surface scheme experiences systematic errors in the atmospheric forcing, or when some physical processes are not properly described. Two conclusions were drawn: soil prognostic variables need to be initialised and short-range forecast errors of near-surface parameters contain information on the state of the soil.

A simple soil analysis scheme (a nudging scheme) was implemented in December 1994. Soil-moisture increments $\Delta\Theta$ at each analysis cycle were assumed to be proportional to increments of specific humidity Δq at the lowest model level produced by the atmospheric analysis:

$$\Delta\Theta = C_{veg} D \Delta q \ ,$$

where D is a global constant coefficient and C_{veg} is the vegetation fraction (introduced in order to reduce correction over deserts). This nudging scheme has proved to be robust enough to control soil-moisture drifts in the root zone during summer. However, recent comparisons of the

ECMWF reanalysis (ERA-15) with observations from the field experiment FIFE (*Betts et al., 1998a*) and the Arkansas Red River basin (*Betts et al., 1998b*) indicate that there are some deficiencies of the nudging scheme. Although the simulated diurnal and annual cycles of soil moisture are both reasonable in ERA-15, they show systematic biases that are compensated by the nudging. The low-level specific humidity has a too strong morning peak and a too low late afternoon minimum, leading successively, to negative and positive increments. Over the Arkansas Red River basin, the nudging exhibits a strong seasonal cycle, with negative values in winter and positive values in summer. These problems are related to a misrepresentation of the vegetation properties in the model (e.g. use of globally constant fields) and to systematic model errors in the radiative and precipitation forcing.

The revised soil analysis scheme

Mahfouf (1991) proposed a simple soil analysis scheme in which soil moisture corrections are expressed as linear combinations of near-surface temperature and relative humidity increments. However, this method could not be implemented because no screen-level analysis was performed at the time at ECMWF.

Given increments of temperature ΔT and relative humidity ΔRH resulting from a two-metre analysis, soil-moisture corrections $\Delta\Theta$ are evaluated as :

$$\Delta\Theta = \alpha\Delta T + \beta\Delta RH$$

The coefficients α and β are computed from optimum interpolation (OI) theory. They explicitly include information on statistics of forecast and observation errors. The statistics of forecast errors were produced from an ensemble of 24-hour simulations of the single-column ECMWF model with different initial soil-moisture conditions. One important feature of the OI scheme with respect to the nudging is the use of more selective criteria of applicability of soil-moisture corrections. Indeed, forecast errors of near-surface parameters are not always informative about soil moisture. As a result, from the statistics of forecast errors, important corrections are only applied with the OI scheme when the temperature and humidity increments have opposite signs (i.e., when the soil is too dry, the boundary layer tends to become too warm and too dry). Optimum coefficients are reduced by an empirical function Φ depending upon environmental conditions when the coupling between the soil and the boundary layer is weak (night-time, winter season, cloudy and rainy conditions, strong advection).

A soil-temperature analysis for the top soil layer T_s (representative of the first 7 cm) is also included. This analysis helps to reduce the forecast errors in the two-metre temperature that are not associated with soil moisture. A simple approach where temperature increments ΔT_s are only expressed as a fraction of the two-metre temperature errors ΔT has been adopted:

$$\Delta T_s = (1-\Phi)\Delta T$$

The empirical function Φ, defined for soil-moisture analysis, is such that temperature analysis is more effective during night and winter, when the two-metre temperature errors are less likely to be related to soil moisture. This way, the two-metre temperature errors are not used to correct both the soil moisture and the soil temperature at the same time.

Screen-level analysis

An analysis of the two-metre temperature and relative humidity has been developed as part of the ECMWF data assimilation system. This analysis is performed every six hours, independently of the atmospheric analysis. The screen-level analysis does not have a direct impact on the atmospheric analysis because the atmospheric model does not have prognostic variables at screen level. However, the screen-level analysis increments are used for the OI soil analysis and, therefore, the atmospheric model feels the influence of the screen-level analysis through the soil variables.

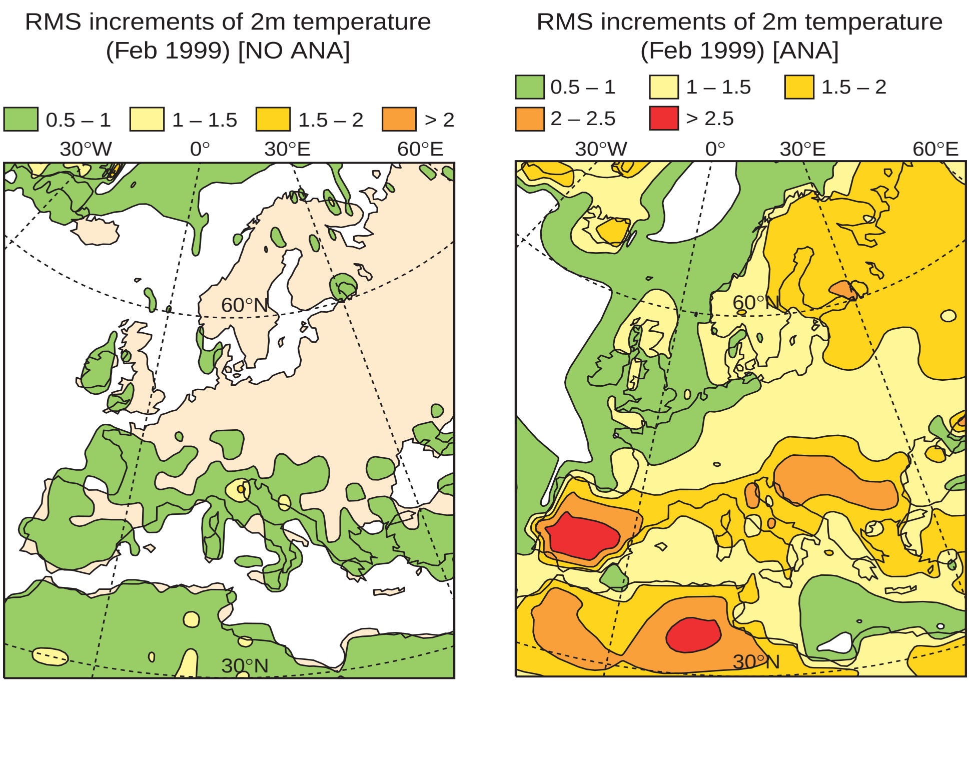

Figure 1 Monthly root-mean-square increments of the two-metre temperature in February 1999 (a) obtained from the assimilation system with upper-air analysis only and (b) including the screen-level analysis.

The screen-level analysis is based on a univariate OI with isotropic structure functions having a correlation length of 300 km. This large radius of influence implies that only large-scale coherent patterns of forecast errors are described by the analysis. Indeed, surface weather parameters are influenced by local effects that cannot be resolved by a global model with 60-km grid. However, in data-sparse areas, the influence of isolated stations can be overestimated. The humidity analysis is performed in terms of relative humidity. This allows an easier specification of forecast errors (the use of specific humidity would have required forecast errors to depend upon the model state) and reduces the influence of mismatches between the model orography and the station heights. The two-metre temperature and relative humidity can be used for other purposes than the soil analysis. These fields are indeed a better representation of the truth because SYNOP measurements are explicitly used to estimate them. This is of particular interest for various applications requiring accurate surface analyses of weather parameters (e.g. forecast verifications, hydrological modelling).

As an illustration of the behaviour of the screen-level analysis, the root-mean-square (RMS) increments of the two-metre temperature over Europe in February 1999, with and without screen-level analyses, are compared in Figure 1 (corresponding to the 'e-suite' period for Cy19r2). Almost negligible corrections (around 0.5 K) are present when only upper-level atmospheric temperatures can modify the surface. The screen-level analysis produces significant RMS increments of between 2 and 3 K over Scandinavia and around the Mediterranean Sea. Corrections at high latitudes are associated with excessive cloud cover over snow-covered areas leading to an overestimation of night-time temperatures. Over southern Europe, the model first-guess is generally too cold. This problem is related to the description of turbulent transfers in stable boundary layers that has already been partially addressed by recent improvements to the land-surface scheme proposed by *Beljaars et al. (1996)*. Even though the magnitude of the RMS increments is much larger when the screen-level analysis is activated, there is some similarity in the patterns, since the largest values are present around the Mediterranean Sea in both cases.

Comparison of the two-metre temperature analysis with SYNOP observations over Europe for a one-week period in February 1999 reveals that, as expected, the analysis field is closer to observations when a screen-level analysis is performed (Figure 2). Large errors (more than 7 K) over Scandinavia can be only slightly reduced by the surface analysis because the observation and background errors in the OI scheme have been set to 2 K and 1.5 K, respectively. Therefore, some of these observations are rejected by the quality control of the analysis. Due to local effects, coastal stations show large departures that are marginally corrected by the analysis when their signal is not consistent with departures from inland stations. The reduction of the standard deviation of analysis departures from 2.91 K to 2.23 K shows that the analysis scheme does not correct for systematic errors only.

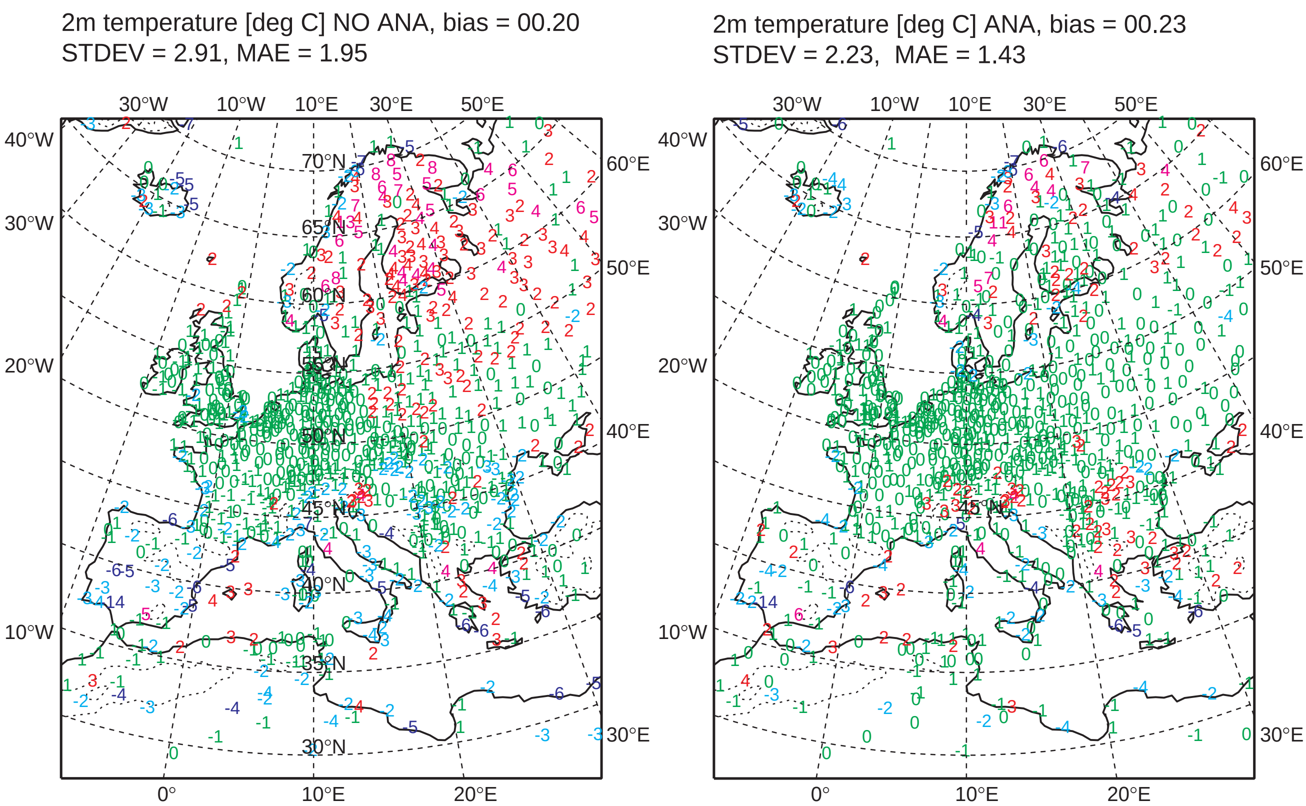

Figure 2 Comparison of two-metre temperature analysis at 00 UTC against SYNOP observations from 4-10 February 1999 for the assimilation system (a) with the upper-air analysis only and (b) including the screen-level analysis.

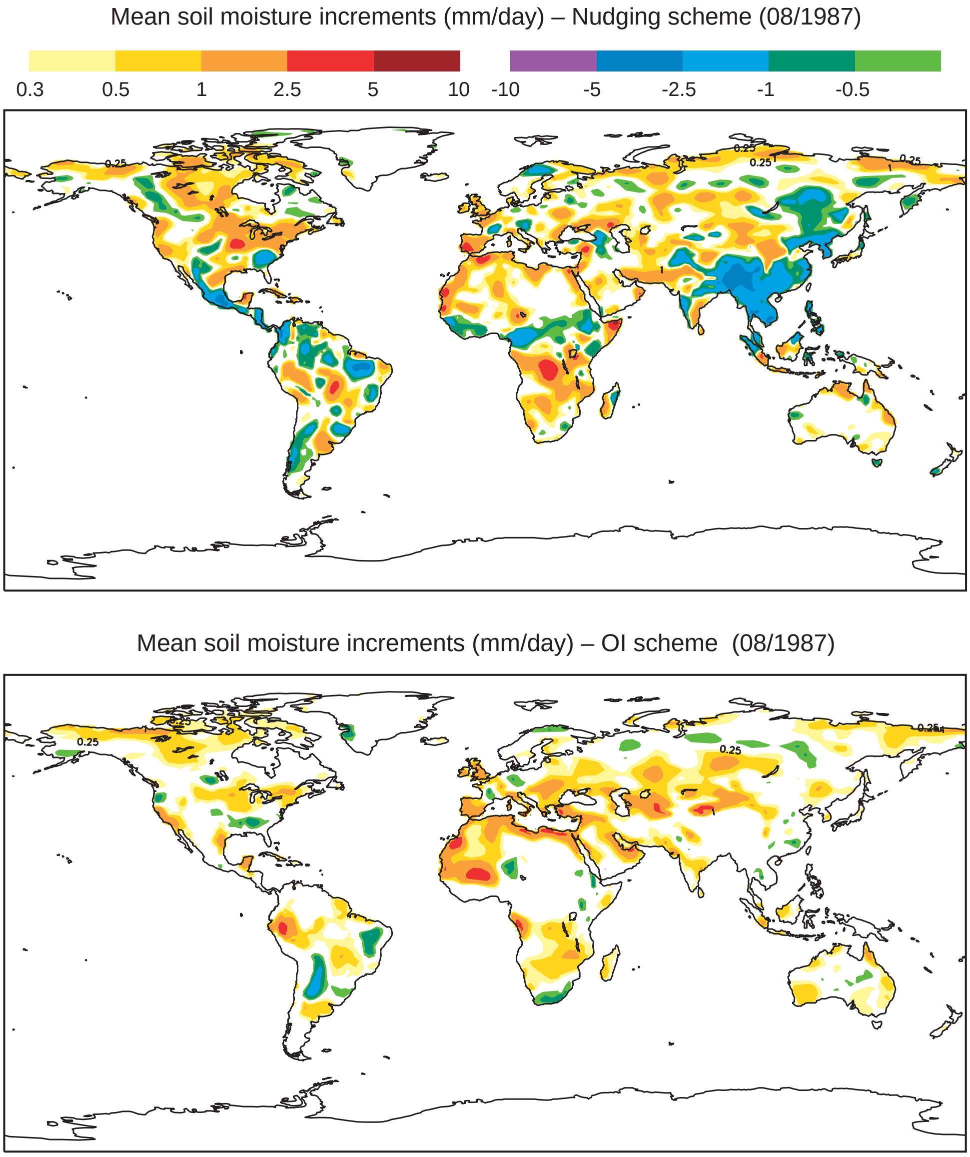

Figure 3 Monthly-mean increments of soil moisture within the root zone (in mm/day) during August 1987 produced by the nudging scheme (upper panel) and by the OI scheme (lower panel).

Tests of the revised soil analysis scheme

In preparation for ERA-40, a series of 3D-Var assimilations at T_L159L60 resolution over a one-year period (July 1987 to June 1988) was performed to compare the two soil analysis schemes. The impact of soil-moisture analysis is more pronounced in summer whereas the impact of soil-temperature analysis dominates in winter.

Monthly-mean increments of soil moisture for August 1987 are presented in Figure 3. There is a consistent moistening over the continents with the two schemes. The drying over Siberia and Canada are also a feature of both analyses. However, corrections in the soil are generally smaller with the OI scheme than with the nudging scheme. This is particularly evident over North America and southern Africa. Over tropical forests, characterised by large negative increments with the nudging scheme, no significant corrections are applied with the OI scheme. Because of the predominance of cloudy and rainy conditions, the soil moisture analysis switches off with the OI scheme. *Douville et al. (1998)* showed that evaporation is usually underestimated over tropical forests because of an overestimation of the canopy resistance in the land-surface scheme. Reducing soil moisture with the nudging scheme has a tendency to further reduce evaporation. During the rainy season soil moisture is close to saturation, which means that evaporation is not controlled by the soil-moisture deficit and, consequently, corrections should be minimal.

The behaviour of the two schemes is examined in more detail over Central US where *Betts et al. (1998a,b)* reported some weaknesses of the nudging scheme used in ERA-15. A 10-day time-series of soil-moisture increments produced by the two analysis schemes for August 1987 is presented in Figure 4. At each analysis cycle, soil-moisture corrections are larger by a factor 2 or 3 with the nudging scheme, and are characterised by a strong diurnal cycle. There is compensation in the nudging scheme between negative increments at the beginning of the day and positive increments at the end of the day. The filtering of the soil-moisture increments by the diurnal cycle appears clearly with the OI scheme where significant corrections are only applied once a day. Monthly-accumulated corrections are much smaller with the OI scheme, which produces a 10 mm increase over

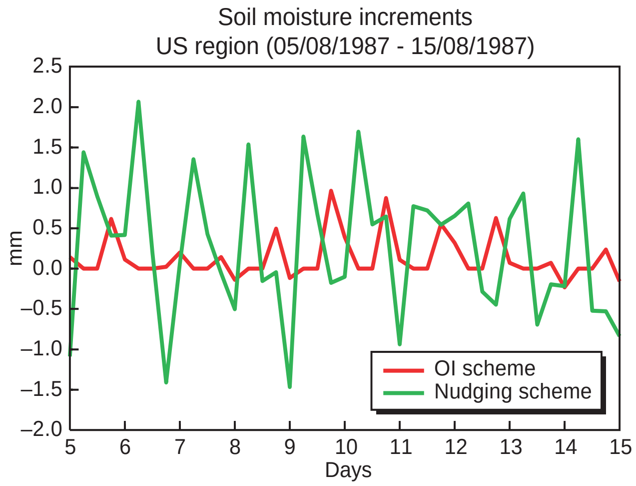

Figure 4 Time evolution of soil moisture increments over the USA during a ten-day period (5-15 August 1987) produced by the nudging scheme and by the OI scheme.

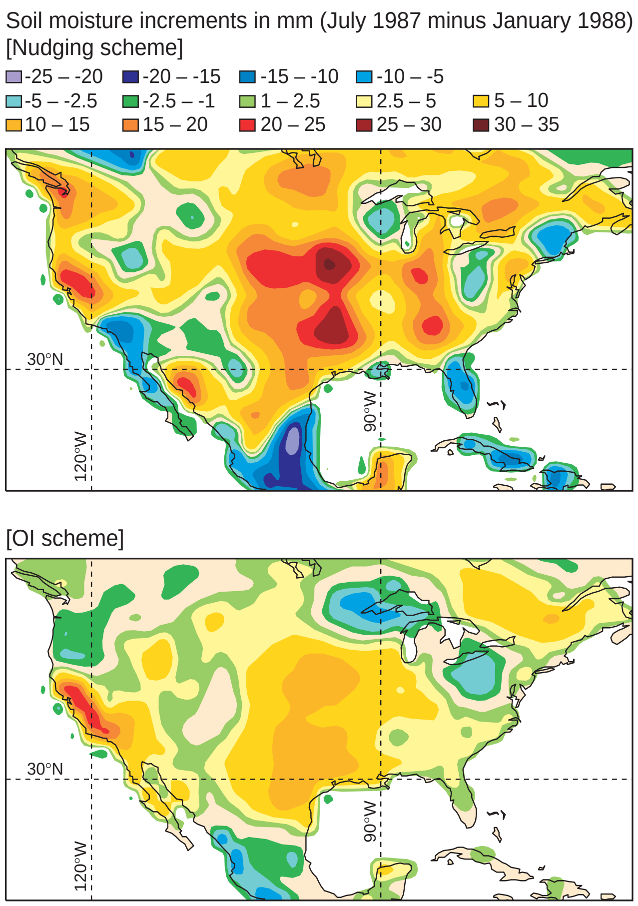

Figure 5 Difference in monthly accumulated soil moisture increments between July 1987 and January 1988 produced by the nudging scheme (upper panel) and by the OI scheme (lower panel).

this region compared with a 30 mm increase with the nudging scheme. The seasonal cycles of soil-moisture increments over the US continent are compared in Figure 5. The monthly differences between July 1987 and January 1988 are mainly positive because both analysis schemes add water in the soil during summer and remove water from the soil during winter. The amplitude of the seasonal cycle is much smaller with the OI scheme since, as shown before, positive soil-moisture increments are reduced in summer, and the increments are minimal in winter because the radiative forcing is low and corrections are strongly reduced through the empirical function Φ.

The soil-temperature analysis is mostly effective in winter at high latitudes. For example, over Siberia during winter the large positive bias in the two-metre temperature induces a significant cooling of the soil. The soil temperature annual cycle over this area for the first soil layer (where the soil corrections are applied), and for the deepest soil layer, is described in Figure 6. When the solar forcing at the surface is non-negligible, soil-temperature corrections are rapidly lost during the short-range forecasts due to the low thermal inertia of the first soil layer. Therefore, the impact of soil-temperature analysis only shows up from November to February in the superficial layer. Despite no explicit corrections, a significant cooling (up to 3 K) of the fourth soil layer is noticed. This effect is produced by soil heat transfer from the surface, and implies a time lag of about three months.

The impact of the revised land-surface analysis is neutral in terms of objective scores for altitude variables, such as the geopotential at 500 hPa. The impact on weather parameters is summarised over Europe in Table 1 for the two-metre temperature and specific humidity during two 14-day periods in June and November 1998, expressed in terms of the mean bias and standard deviation with respect to SYNOP observations for the 72-hour range. The OI scheme systematically produces lower values of the standard deviation of forecast errors than the nudging scheme. The bias in specific humidity is slightly reduced. The reduction of the positive bias in winter is a result of the soil-temperature analysis.

Parameter	Bias (NUDG)	Stdev (NUDG)	Bias (OI)	Stdev (OI)
T (November) [K]	+0.68	2.51	+0.63	2.45
Q (November) [g/kg]	+0.28	0.87	+0.25	0.87
T (June) [K]	+0.12	3.18	+0.13	3.10
Q (June) [g/kg]	–0.87	2.24	–0.85	2.10

Table 1 Two-week averages of 72-hour forecast errors of the two-metre temperature and specific humidity over Europe in November 1998 and June 1998 from 3D-Var experiments with a soil analysis scheme based on nudging (NUDG) and with a soil analysis scheme based on optimum interpolation (OI)

Summary of the impact of the new scheme

Two important changes on the land-surface analysis implemented in 1999 have been explained and illustrated in this paper. The analyses of two-metre temperature and relative humidity produced by the ECMWF data assimilation system since March 1999 are of better quality because they explicitly use SYNOP observations through a simple univariate OI. These fields are, therefore, less contaminated by model biases and can be used with more confidence for various meteorological applications. The two-metre analysis increments are used as input for a revised soil analysis scheme based on OI for both the water content in the root zone and for the temperature. Results show that the revised scheme has a better overall performance than the nudging scheme. The compensating effect of the diurnal cycle of soil-moisture increments is strongly reduced with the OI scheme. Monthly-mean increments in most regions of the globe are of comparable size or smaller with the OI scheme. The damping of the seasonal cycle observed in ERA-15 is also reduced. In terms of forecast impact, there is a small but systematic reduction of the forecast errors on weather parameters. The effect of the soil-temperature analysis is to reduce significantly the warm bias over snow-covered areas at high latitudes in winter.

The revised soil analysis scheme (like the previous scheme) still compensates for systematic errors coming from an inaccurate representation of some physical processes in the land-surface scheme. Preliminary experiments with an improved land-surface scheme and new climatological datasets show much smaller soil-moisture increments, thereby reducing this undesirable feature of the soil analysis. The OI soil analysis was put into operational use in July 1999 and will be used for ERA-40 with the revised land-surface scheme. It is hoped to produce an improved soil and surface climate from this reanalysis project, an objective that is particularly important, given the lack of direct observations on these variables.

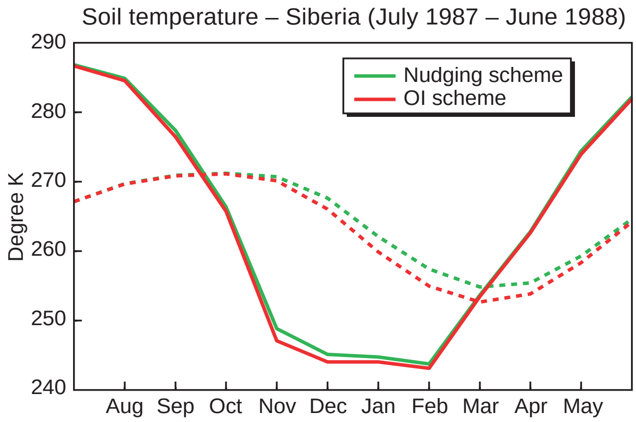

Figure 6 Annual cycle of monthly-mean values of analysed soil temperature over Siberia from July 1987 to June 1988. The solid lines correspond to the first soil layer (7 cm) and the dashed lines correspond to the fourth soil layer (1.98 cm). In the OI scheme a soil-temperature analysis of the first soil layer is performed, whereas there is no soil-temperature analysis with the nudging scheme.

References

Beljaars, A., J.-F. Mahfouf, J. Teixeira and **P. Viterbo**, 1996: Improvements to the 2m temperature forecasts. *ECMWF Newsletter*, 73, 2-6.

Betts, A.K., P. Viterbo, and **A.C.M. Beljaars**, 1998a: Comparison of the land surface interaction in the ECMWF reanalysis model with the 1987 FIFE data. *Monthly Weather Review*, **126**, 186-198.

Betts, A.K., P. Viterbo and **E. Wood**, 1998b: Surface energy and water balance for the Arkansas Red River basin from the ECMWF reanalysis. *J. Climate*, **11**, 2881-2897

Douville, H., J.-F. Mahfouf, S. Saarinen and **P. Viterbo**, 1998: The ECMWF surface analysis: diagnostics and prospects. *ECMWF Technical Memorandum No 258*, 51 pp.

Mahfouf, J.-F., 1991: Analysis of soil moisture from near-surface parameters: A feasibility study. *Journal of Applied Meteorology*, **30**, 1534-1547

Viterbo, P., 1995: Initial values of soil water and the quality of summer forecasts. *ECMWF Newsletter*, **69**, 2-8

Viterbo, P. and **A.C.M. Beljaars**, 1995: A new land surface parametrization scheme in the ECMWF model and its validation. *ECMWF Technical Report No 76*

J.-F. Mahfouf, P. Viterbo, H. Douville, A.C.M. Beljaars and S. Saarinen

IFS tests using MPI/OpenMP

OpenMP is a relatively new standard for parallel programming on shared memory multiprocessor systems. This article provides some background on OpenMP and how it is being applied in the IFS.

The IFS has been running operationally for several years on the VPP700 and, more recently, on the VPP5000. Both these systems are **distributed memory** supercomputers and are programmed using Fortran 90 for the scientific code and the Message Passing Interface (MPI) standard for handling communication between processors. During the past few years, many of the major computer manufacturers have developed Symmetric Multi-Processor (SMP) systems, which provide a **shared memory** programming model to the user even though the underlying memory packaging may be distributed. An example of such a system is the SGI Origin 2000 (e.g. STYX and HADES).

A major advantage of shared memory systems over distributed memory systems is their ease of programming. With shared memory an incremental approach to parallelisation can be pursued - a time consuming section of code can be parallelised and correct execution confirmed before proceeding on to the next code section. This approach is not possible with distributed memory programming where issues of data distribution, message passing and synchronisation impose a heavy burden on the programmer.

In the past, shared memory machines such as the CRAY C90 were programmed using macrotasking (as used at ECMWF) or microtasking. However, each manufacturer had their own library calls and directives to express parallelism. Attempts to provide a standard for shared memory programming were pursued by the Parallel Computing Forum (an informal industry group) and later by the ANSI/X3 authorised subcommittee X3H5. Draft standards were written, but these were either incomplete or lacked the functionality of the manufacturer-specific implementations.

This has changed in recent years with the availability of **OpenMP**. Now applications can be parallelised for SMPs and, most importantly, remain portable on such systems. However, if a single application is required to run on more than one SMP, MPI is needed. This can be done by running MPI processes on each processor of each SMP or, alternatively, by running only one MPI process per SMP and OpenMP within each SMP. The latter technique could be particularly attractive for applications where the cost of global communication is high on large numbers of processors. This is shown graphically in Figure 1 where explicit (MPI only) communications paths are shown for a global operation (all-to-all) on two 4-processor SMP nodes. In contrast, Figure 2 shows an MPI/OpenMP approach where MPI communication only occurs once between SMP nodes. This approach is expected to result in a substantial reduction in explicit MPI communication and, thereby, in better performance on large SMP systems compared with an MPI-only approach. Note that SMP nodes in a few years time will probably have 32 to 128 processors, and so systems having 10 to 100 of such nodes would not be unrealistic for NWP centres.

In early 1999 a project was started at ECMWF to obtain a greater understanding of OpenMP and its applicability to the IFS. After experimenting with some kernel tests it was decided to add OpenMP directives to port the IFS forecast model. The initial version took approximately 2-3 months to complete and was included in IFS Clearcase cycle CY21r3. This was subsequently optimised for cycles CY21r5 and

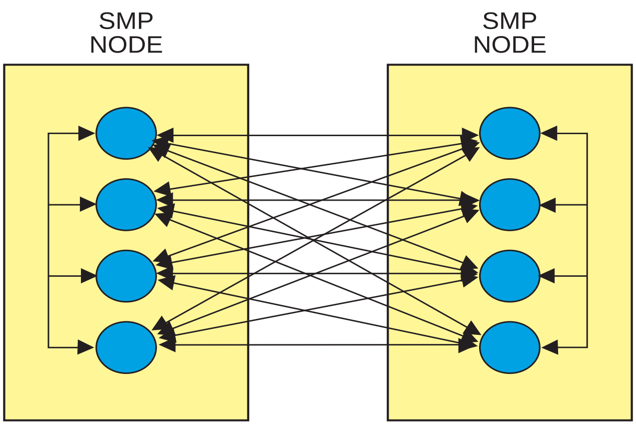

Figure 1 Explicit MPI global communication for two four-processor SMP nodes using an MPI-only approach.

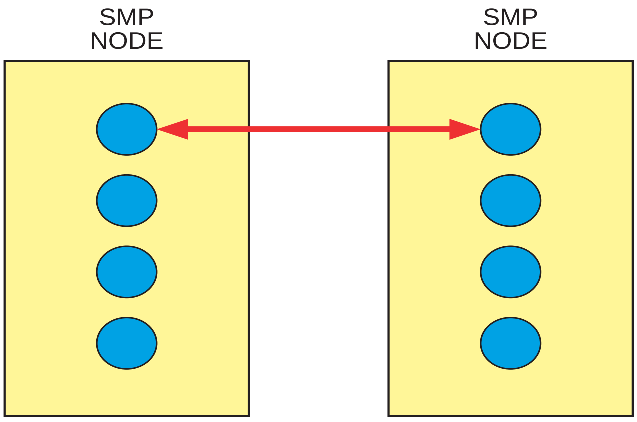

Figure 2 Explicit MPI global communication for two 4-processor SMP nodes using an MPI/OpenMP approach (one MPI process/ four OpenMP threads per node).

CY22r1 based on experience gained from tests on a number of SMP systems.

On a multi-node system, the IFS implementation permits both MPI and OpenMP parallelism to be exploited. The only constraint is that the number of OpenMP threads per node should not exceed the number of processors on a node. If only MPI parallelism is required (e.g. for a distributed memory system) then the same source can be used without change.

In total 30 IFS routines were modified to support OpenMP for the forecast model. Figure 3 shows an extract from scan2mdm.F90. OpenMP directives appear in the source code as comments starting with the sentinel !$OMP. In this extract the SCHEDULE clause determines how DO loop iterants are to be assigned to processors and, in this case, each iteration (NPROMA stride) is assigned dynamically. By this simple mechanism the effect of computational imbalance arising in the Centre's physics package can be minimised. Of course, this is constrained by the static distribution of grid points to MPI processes in the IFS initialisation.

Tools to analyse OpenMP performance (*Guide/Guideview*) and check the correctness of OpenMP constructs (*Assure/Assureview*) are under evaluation at ECMWF (Figure 4). These tools are expected to help programmers diagnose the performance of each OpenMP construct and help identify sections of code where OpenMP should be applied.

At present the IFS MPI/OpenMP implementation has been verified on the following SMP systems,

- FUJITSU E10000
- IBM (Nighthawk, Winterhawk),
- NEC SX4
- SGI Origin 2000

On each of these **single node** systems the MPI-only approach produced better performance over a combined MPI/OpenMP approach, with the NEC SX4 showing the least difference. Future improvements to the IFS OpenMP implementation are expected bring OpenMP performance even closer to MPI performance on these systems. The MPI/OpenMP performance is shown in Figure 5 for an SGI Origin 2000 and in Figure 6 for an NEC SX/4.

```
!$OMP PARALLEL PRIVATE(JKGLO,ICEND,IBL,IOFF,ZGT1A,ZWORK)
        IF (.NOT.ALLOCATED(ZGT1A)) ALLOCATE(ZGT1A(NPROMA,IFLDSGT1A))
!$OMP DO SCHEDULE(DYNAMIC,1)
        DO JKGLO=1,NGPTOT,NPROMA
          ICEND=MIN(NPROMA,NGPTOT-JKGLO+1)
          IBL=(JKGLO-1)/NPROMA+1
          IOFF=JKGLO
          CALL CPGLAG(CDCONF(4:4),ISTROWL,IENDROWL,ICEND,JKGLO &
           &,ICEND,KWORK,ZWORK,IMLATLAG,NASLB1,ZSLBUF1 &
           &,GPP(1,1,IBL),RCORI(IOFF),GEMU(IOFF)&
           &,GELAM(IOFF),GESLO(IOFF),GECLO(IOFF),GM(IOFF)&
           &,OROG(IOFF),GNORDL(IOFF),GNORDM(IOFF)&
           &,GSQM2(IOFF),RCOLON(IOFF),RSILON(IOFF)&
           &,RINDX(IOFF),RINDY(IOFF),GAW(IOFF)&
           &,VMAXTORO(IOFF)&
           &,RSTOPHU(IOFF),RSTOPHV(IOFF)&
           &,RSTOPHT(IOFF),RSTOPHQ(IOFF)&
           &,ZGT1A(1,IUT1) ,ZGT1A(1,IVT1) ,ZGT1A(1,ITT1)&
           &,ZGT1A(1,IQT1) ,ZGT1A(1,IO3T1),ZGT1A(1,ILT1)&
           &,ZGT1A(1,IIT1) ,ZGT1A(1,ISPD1),ZGT1A(1,ISVD1)&
           &,ZGT1A(1,ISVT1),ZGT1A(1,ISPT1),ZGT1A(1,IAT1)&
           &,ZGT1A(1,IVVELPP), ZGT1A(1,ITPP))

          GPP(1:ICEND,1:IFLDSGT1A,IBL)=ZGT1A(1:ICEND,1:IFLDSGT1A)
        ENDDO
!$OMP END DO
        IF (ALLOCATED(ZGT1A)) DEALLOCATE (ZGT1A)
!$OMP END PARALLEL
```

Figure 3 OpenMP code extract from control/scan2mdm.F90 (call to cpglag).

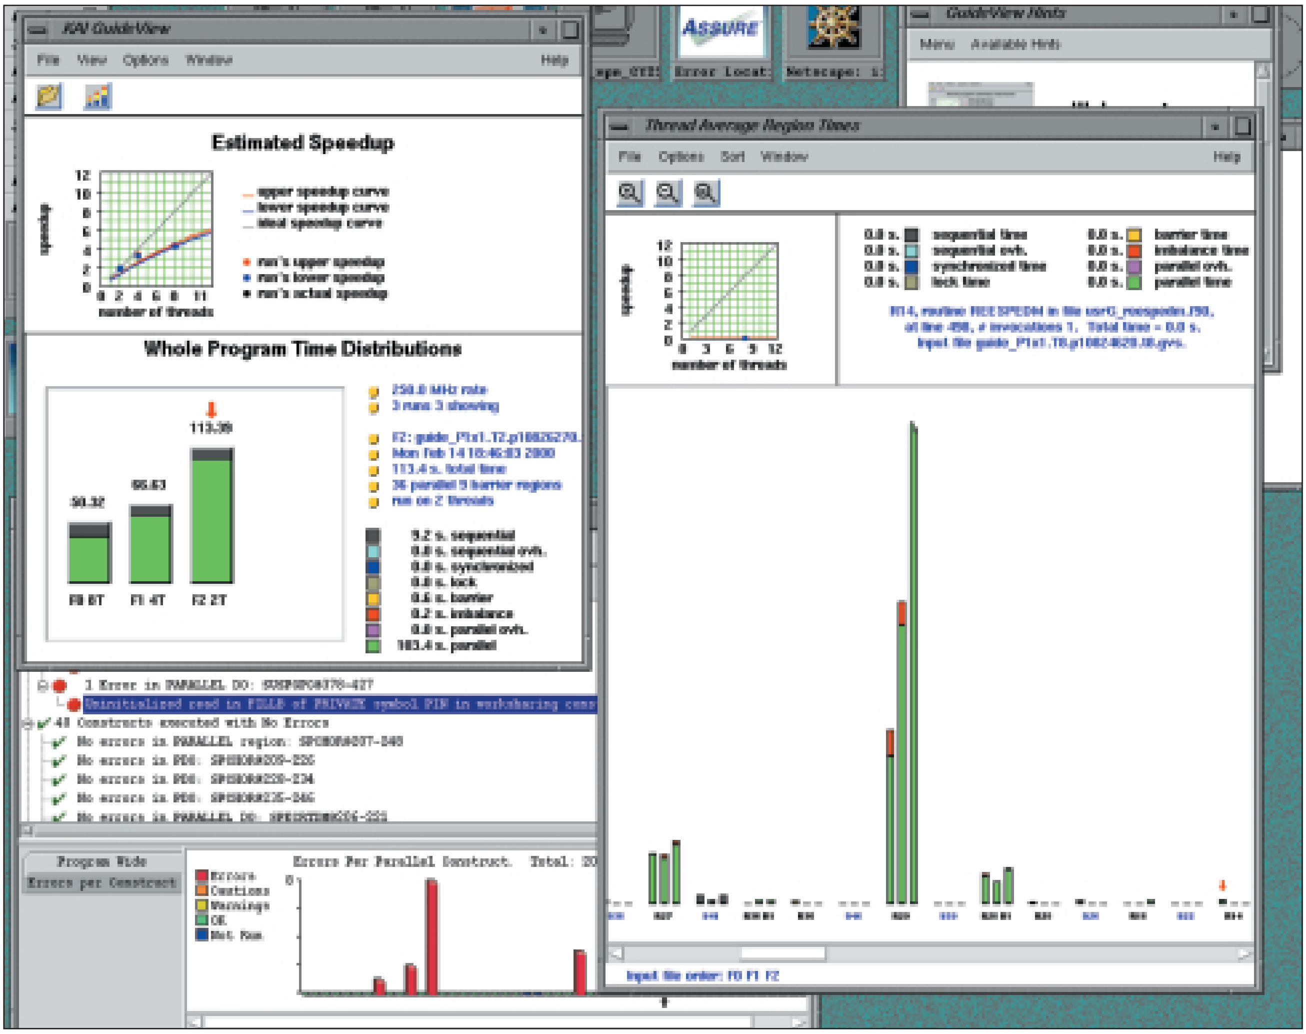

Figure 4 OpenMP tools Guideview/Assureview

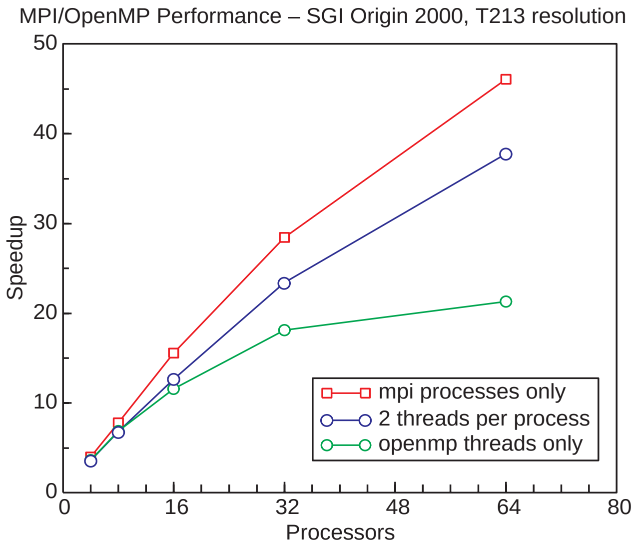

Figure 5 IFS T213 forecast model performance on an SGI Origin 2000 using different combinations of MPI and OpenMP.

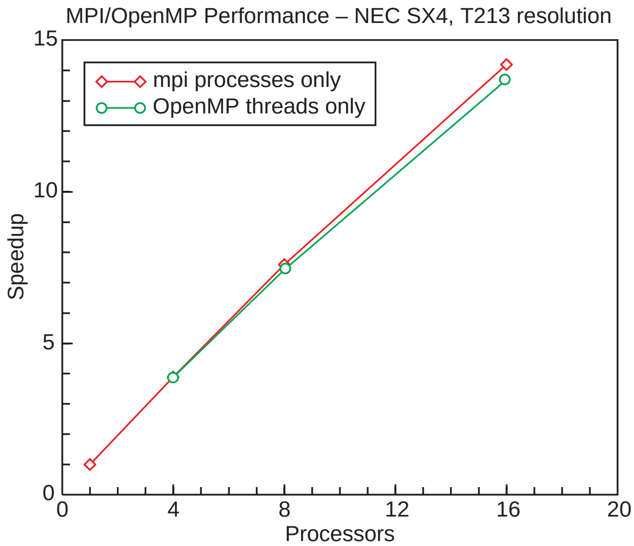

Figure 6 IFS T213 forecast model performance on an NEC SX/4 comparing MPI and OpenMP.

However, it should be noted that the target architecture for using MPI/OpenMP is a **multiple node** system. On such systems it is expected that the use of OpenMP on each node and MPI between nodes will provide better performance than using just MPI over all processors/nodes.

Time will tell!

Glossary

MPI Message Passing Interface, a library for parallel programming on distributed memory systems (such as the Centre's Fujitsu VPP700 and VPP5000 systems). MPI can also be used on shared memory systems. See http://www.mpi-forum.org/

OpenMP A directive-based standard for parallel programming on shared memory systems. See http://www.openmp.org/

Process This is a general term for a program that is being executed. All work done by a processor (CPU) contributes to the execution of processes. When an IFS experiment is running on several Fujitsu VPP processors it has one **MPI process** per processor.

SMP Symmetric Multi-Processor, a computer system with a number of processors each having access to a single shared memory, e.g. *styx* and *hades* are SMPs each with 16 processors.

Thread This is different from a process or task. A thread (sometimes called a lightweight process) is a kind of stripped-down process – it is just one 'active hand' in a program – something that the CPU is doing on behalf of a program, but not enough to be called a complete process. Threads remember what they have done separately, but they share the information about which resources a program is using, and what state the program is in. A thread is only a CPU assignment. **OpenMP** uses **threads** to distribute work among processors on an SMP.

George Mozdzynski

ECMWF publications

Meteorological Bulletins

M1.5/2 IFS Documentation Part I: Observation Processing (CY21r4), edited by Peter W White, *March 2000*

M1.5/3 IFS Docmentation Part II: Data Assimilation (CY21r4), edited by Peter W White, *February 2000*

M1.6/4 IFS Documentation Part III: Dynamics and Numerical Procedures (CY21r4), edited by Peter W White, *March 2000*

M1.6/6 IFS Docmentation Part V: The Ensemble Prediction System (CY21r4), edited by Peter W White, *February 2000*

Technical Memoranda

301 **M. Janisková**, **J-F. Mahfouf**, **J-J. Morcrette** and **F. Chevallier**: Development of linearized radiation and cloud schemes for the assimilation of cloud properties. *March 2000*

302 **V. Marécal**, **E. Gérard**, **J-F. Mahfouf** and **P. Bauer**: The comparative impact of the assimilation of SSM/I and TMI brightness temperatures in the ECMWF 4D-Var system. *May 2000*

303 **J. Segschneider**, **D.L.T. Anderson**, and **T.N. Stockdale**: Towards the use of altimetry for operational seasonal forecasting. *June 2000*

304 **J. Segschneider**, **M. Balmaseda**, and **D.L.T. Anderson**: Anomalous temperature and salinity variations in the tropical Atlantic: possible causes and implications for the use of altimeter data. *June 2000*

305 **A. Troccoli**, **M. Balmaseda**, **J. Segschneider**, **J. Vialard**, **D.L.T. Anderson**, **K. Haines**, **T.N. Stockdale** and **F. Vitart**: Salinity adjustments in the presence of temperature data assimilation. *June 2000*

306 **A.T. Weaver** and **P. Courtier**: Correlation modelling on the sphere using a generalized diffusion equation. *June 2000*

307 **S.L. Mullen** and **R. Buizza**: Quantitative precipitation forecasts over the United States by the ECMWF Ensemble Prediction System. *June 2000*

ERA Report

ERA40 PRS1 **A.J. Simmons** and **J.K. Gibson** The ERA-40 project plan. *July 1999*

EUMETSAT/ECMWF Report

EUMETSAT/ECMWF Numerical Weather Prediction Satellite Application Facility Research Reports No. 2.: **F. Chevallier** and **J-F. Mahfouf**: Evaluation of the Jacobians of infrared radiation models for variational data assimilation, *May 2000*

Workshop Proceedings

ECMWF/EUMETSAT Workshop on "Use of ATOVS data for NWP assimilation". *2–5 November 1999*

Seminar Proceedings

Seminar on "Diagnosis of models and data assimilation systems". *6–10 September 1999*