In this issue

This Newsletter is edited and produced by User Support and is designed and printed by ECMWF.

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, RG2 9AX, UK

Fax: (0118) 986 9450

Telephone: National....0118 949 9000

International....+44 118 949 9000

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

Editorial

Many users of ECMWF forecast products had reported a severe cold bias in the 2m temperature in winter, particularly over 1995/6. The first article describes what was done to identify the problem, and the model changes implemented to reduce the temperature bias.

The production phase of the ECMWF Re-analysis project is drawing to a close. The start of the project was reported three years ago (ECMWF Newsletter No. 64), the article in this issue (page 7) describes what has been achieved since then.

More and more use is being made of ECMWF's Ensemble Prediction System (EPS) to estimate the level of confidence that can be attached to the forecasts. One area of interest is the use of EPS to provide confidence level information in the more extreme weather events that the forecast predicts. A study of three cases of predicted heavy rainfall is given on page 17. Another recent example of severe weather (widespread flooding over western Europe in January 1995) is the topic of a recent ECMWF Technical Report (No. 77). A summary of the report is given on page 24

A new parametrization of sub-grid scale orographic drag was incorporated into the operational forecast last year. An ECMWF Technical Memorandum (No. 218) provides the background to this change, this memorandum is summarised on page 25.

ECMWF has been developing and using parallel processing methods on its computer systems since 1983. Previous Newsletter articles have covered some of the work on shared memory systems, the article in this issue on page 26 reviews the history, current situation and future plans of ECMWF's involvement in distributed memory systems.

Smart card access control to ECMWF's computer systems is being introduced in stages. The reasons for this scheme, and its initial implementation were covered in previous Newsletter articles, the article in this issue on page 30 covers current progress and future plans in this area.

The major new standard for Fortran (known as Fortran 90) is being used more and more at ECMWF. An update, to be known as Fortran 95, is expected to become an ISO standard very shortly. The article on page 31 describes the background, and summarises the changes.

Changes to the operational forecasting system

Recent changes

- A technical error affecting the computation of sub-grid orography processes in the forecast model which was introduced with the change to sea surface temperature and sea ice on 23 April 1996 was corrected on 31 May.
- The complete operational suite was implemented on the Fujitsu VPP700 on 18 September:
  - (i) Main T213 ten-day forecast suite
  - (ii) Optional project wave forecast suite
  - (iii) 00 UTC optional project forecast suite
  - (iv) Ensemble prediction suite.

  Together with the move to the VPP700 a new model cycle (CY15R5) and changes to the 3D-Var code were introduced:
- model cycle CY15R5, introduced on 18 September, had the following changes:
  - a) revision of boundary layer diffusion and introduction of soil moisture freezing. The effect is a reduction of the near-surface temperature errors in stable situations. It implies a reduction of the night time temperature errors over land in summer and a significant reduction of the winter cold bias of day and night time forecasts;
  - b) a revised semi-Lagrangian treatment of the thermodynamic equation leading to smoother meteorological fields over steep orography.

  This cycle also contained changes necessary for the migration to the Fujitsu, none of which have any detectable meteorological impact, at the same time,
- the 3D-Var code has been generalised to allow it to be run on parallel, distributed-memory computers using message passing. A new parallel code is used for observation pre-processing and screening. Quality control and the calculation of background error variances are now based on the variational analysis rather than the earlier OI approach. Objective verification indicates little overall sensitivity of forecast performance to these changes.

Bernard Strauss

Improvements to the 2m temperature forecasts

Introduction

Many users of ECMWF products have reported a severe cold bias in the 2m temperature forecasts in winter. The problem was particularly serious in the winter of 1995/1996, mainly as a result of the anomalous flow circulation with a persistent high pressure area over continental Europe. The standard verification with SYNOP observations has shown that the problem was more pronounced in areas below freezing than in areas above zero, as illustrated in Fig. 1 for Berlin. The latter suggests a possible link between the model bias and soil moisture freezing, a process that is currently not represented in the model. Daily verification with SYNOP data also shows that the night time temperatures are too low in all seasons, and correspondingly the diurnal cycle is too large.

The Research Department has investigated this problem extensively making use of operational and non operational observations. With the help of tower observations covering the lowest two model levels (Cabauw, The Netherlands) it could be shown that for weak winds, a steep inversion develops in the model between the lowest model level (30m above the surface) and the layers above (140m and up). Such shallow inversions, with cold air near the surface and warm air aloft, do occur, but the model tends to produce them more frequently and with too large temperature gradients. Once such an inversion has formed, vertical diffusion by turbulence is suppressed and the lowest model level is completely decoupled from the upper air.

The second result from the diagnostic study is that the soil over continental areas tends to cool excessively. The soil temperature observations as received from some member states have been very useful in diagnosing the nature and magnitude of the problem. An example is given in Fig. 2 for an area in Germany. The comparison of 20 cm deep temperature observations with operational 24 hour forecasts clearly indicates the cold bias in the model, and it also shows the levelling-off of the observations at 0°C, whereas the model soil temperatures do not feel the 0° thermal barrier. It should be remembered that the land surface temperatures in the current scheme are the result of heating and cooling provided by the atmospheric model during the first guess computations (6 hour forecast). The seasonal evolution of the soil temperatures is determined by the model and by the seasonal evolution of the atmosphere during data assimilation, and no soil temperature observations are used directly.

Two conclusions emerged:

(i) The vertical diffusion of heat is too weak in stable situations allowing for too much decoupling between atmospheric and soil temperatures, and

(ii) soil moisture freezing, which was not represented in the model, plays a rather important role in the evolution of soil temperatures in winter. The latter is suggested by the nature of the 2m temperature errors in Fig. 1, but it is also rather obvious from time series of soil temperature observations in Germany (Fig. 2).

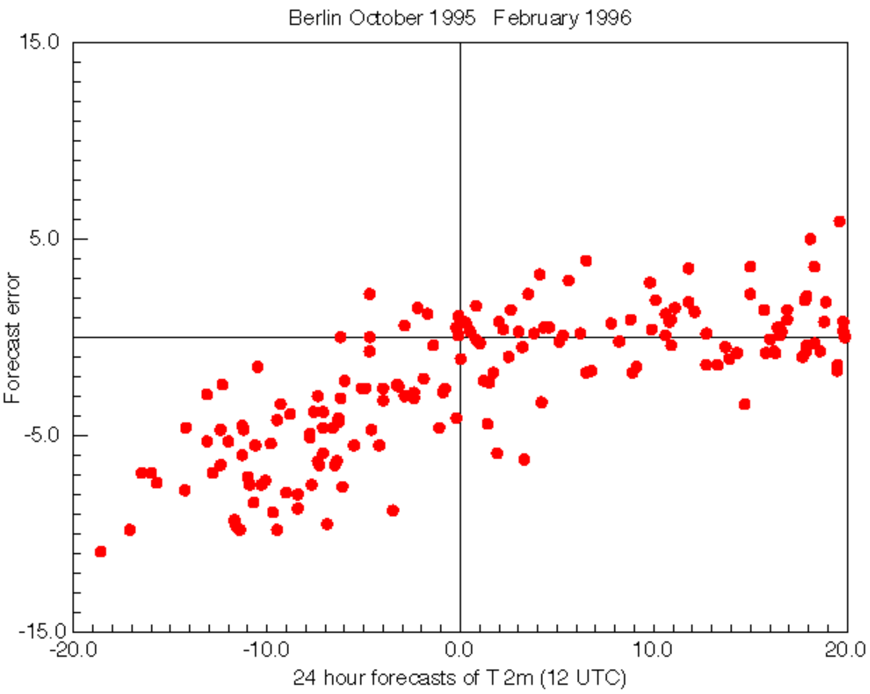

***Fig. 1:** Errors of 24 hour forecasts of 2m temperature at Berlin as a function of temperature.*

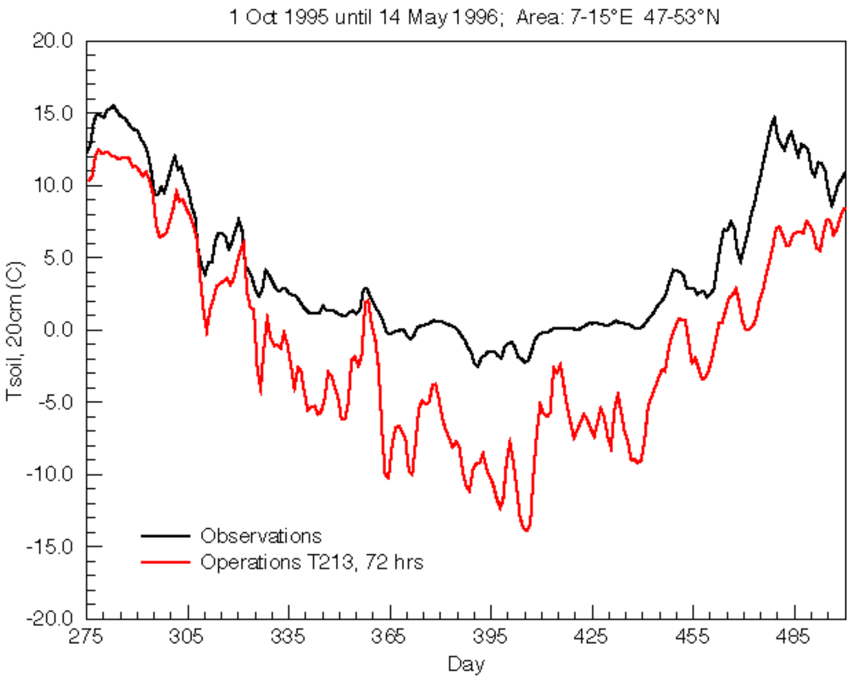

***Fig. 2:** Time series of observed soil temperatures (20 cm deep) compared with 72 hour forecasts of the operational ECMWF model in an area over Germany. The time axis is day number counted from 1-1-95; day 275 corresponds to 1 October 1995 and day 500 corresponds to 14 May 1996. The area includes about 100 SYNOP stations.*

Model changes

In order to improve the model, a package of three changes has been prepared for operational implementation. It consists of:

(i) increased turbulent diffusion of heat in the stable boundary layer,

(ii) introduction of soil freezing, and

(iii) increased skin layer conductivity.

The increased turbulent diffusion spreads the surface cooling over a deeper atmospheric layer leading to less cold surface temperatures, the soil freezing slows the soil temperature drop near freezing, and the higher skin layer conductivity increases the coupling of the radiative surface temperature with the underlying soil and reduces the amplitude of the diurnal temperature cycle. All three changes have a considerable impact on the near surface temperatures, but soil moisture freezing is the dominant effect as far as the seasonal evolution of soil temperatures is concerned. This is illustrated by a simple calculation, which shows that the amount of energy necessary to freeze/thaw 1 m^3 of wet soil, would cool/warm this soil by 49 K if the phase transition was not taken into account.

Effect on seasonal time scales

The changes, discussed above, affect the seasonal evolution of soil temperatures and would need a long data assimilation experiment to test. Such experiments are very demanding on computer resources. As a cheap alter-

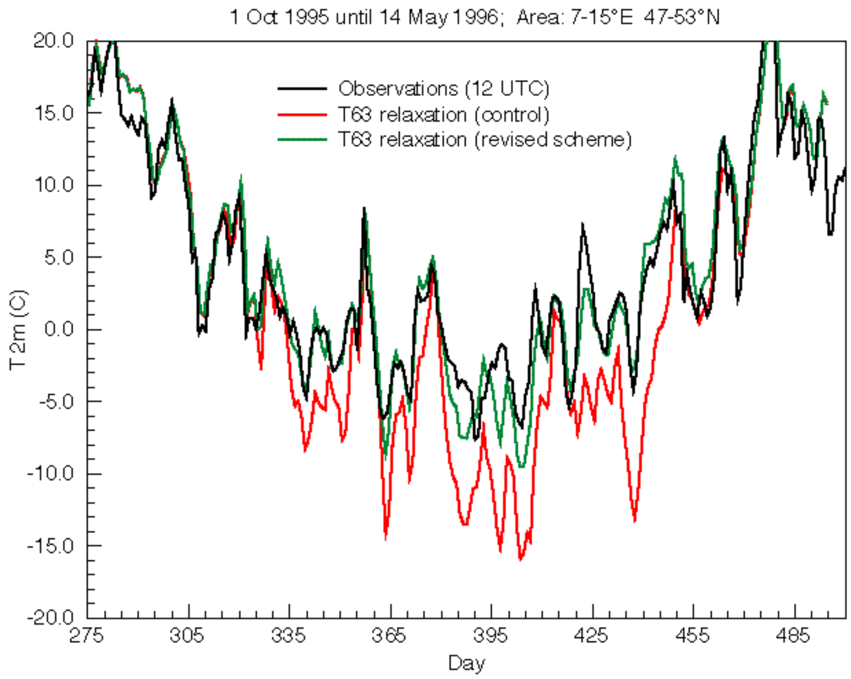

Fig. 3: *Area averages (7-15E 47-53N) of 12z temperatures at 2 m from long runs at T63 resolution with relaxation towards the analysis above the boundary layer compared with observations.*

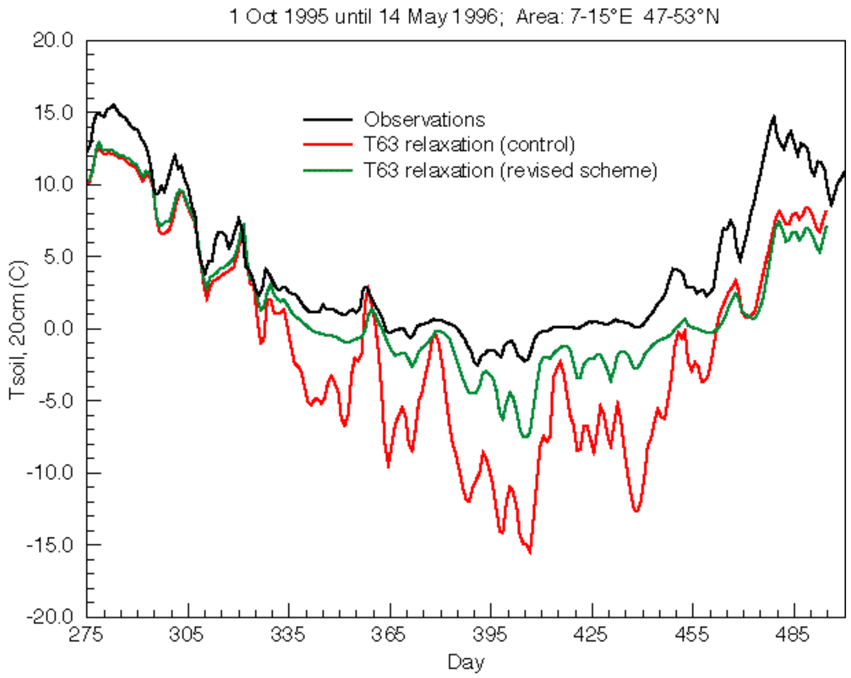

Fig. 4: *Area averages (7-15E 47-53N) of soil temperatures (20 cm deep) from long runs at T63 resolution with relaxation towards the analysis above the boundary layer compared with observations.*

native to data assimilation, long integrations are used at T63 resolution in which the model atmosphere above the boundary layer is relaxed towards the analysis. In this way the synoptic situation of the winter of 1995/1996 could be reproduced realistically with a single long integration. First it was verified that the control run with relaxation towards the analysis, reproduced the cold bias problem and then experiments were carried out with the changes described above. The long experiments at T63 resolution start from 1 October 1995 and finish in May 1996 in order to run through a freezing and melting season.

Results from the long relaxation integrations are shown in Figs. 3 and 4 for an area in Germany. The operational model and the revised model are compared with observed temperatures at screen level and with observed soil temperatures both averaged over the same domain. It is clear that the revised model alleviates the 2m cold bias and that the soil temperature evolution is much closer to the observations. Sensitivity experiments with different parts of the new parametrization have shown that about half of the improvement on the day time 2m temperatures in winter comes from the increased turbulent diffusion; the other half is due to the introduction of freezing. About three quarters of the improvement of deep soil temperatures can be attributed to the effect of soil freezing and the remaining part to increased turbulent diffusion. The skin layer conductivity has no impact on the seasonal temperature

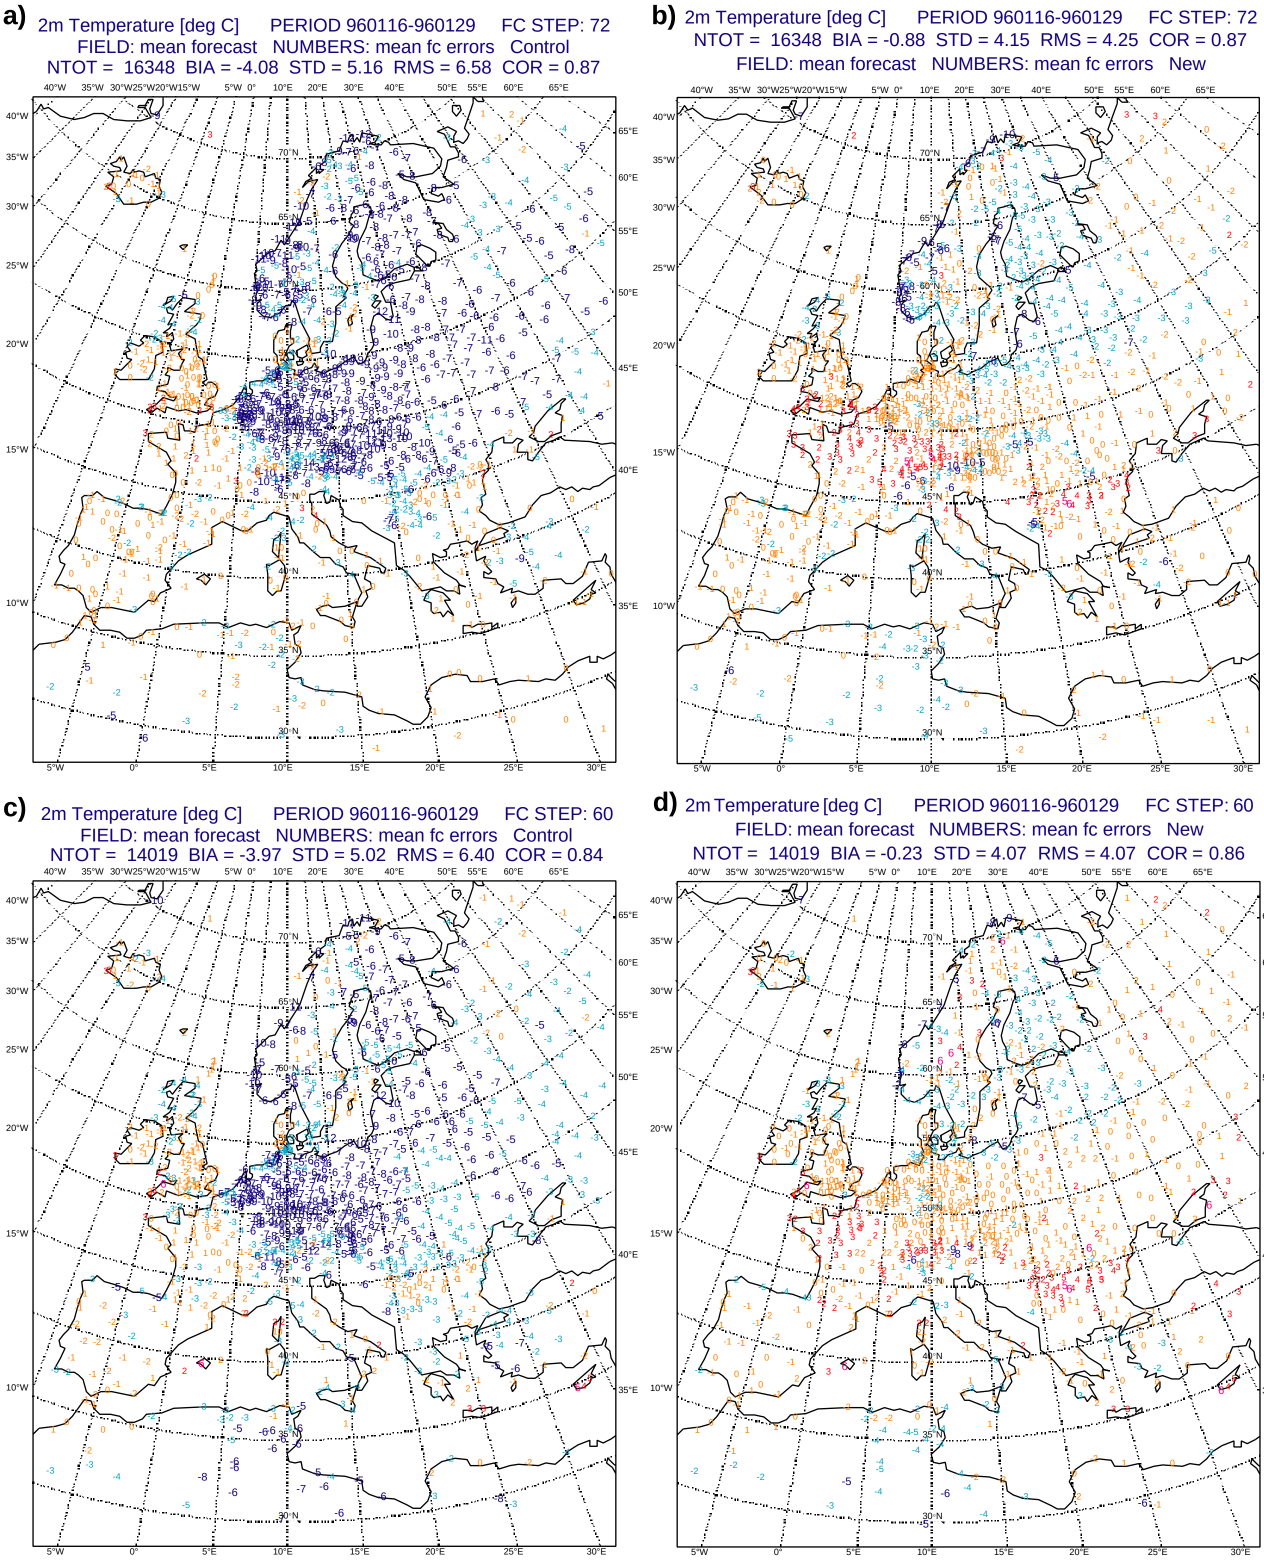

Fig. 5: *Temperature errors at the 2 m level averaged over all 72 (a,b) and 60 (c,d) hour forecasts from the data assimilation experiment with the revised model (Figs. b,d) compared with the control experiment (Figs. a,c). The printed numbers are the mean errors at SYNOP stations for forecasts with initial dates between 16 and 29 January 1996.*

evolution, it only affects the night time temperatures and the amplitude of the diurnal cycle.

Effect on 10 day forecasts

In order to investigate the impact on the diurnal temperature cycle and on the day and night time 2m temperature forecasts, data assimilation/forecast experiments were performed at T106 resolution for January and April 1996.

Both experiments were initialized with soil temperatures from the long relaxation runs.

The 2m temperature errors are shown for SYNOP locations averaged over 14 consecutive 60 and 72 hour forecasts in January with the old and the new scheme respectively (Fig. 5). A large portion of the systematic errors in operations is eliminated with the new scheme. Area averaged biases are reduced from -4.0° to -0.2° at night and from -4.1° to -0.9° for the midday forecasts. The impact on forecast performance of the large scale weather patterns (e.g. expressed by anomaly correlation of 500 and 1000 hPa fields) is small. This is because the near surface temperature biases are limited to a very shallow layer with negligible implications for the large scale height fields or synoptic developments.

April gives a slightly different picture. The daytime temperature bias over central Europe has already disappeared because of the solar heating. The improvement is mainly in the night time temperatures and in the amplitude of the diurnal cycle. This is illustrated by time series of consecutive 54, 60, 66 and 72 hour forecasts for South and Northern Europe, averaged over all SYNOP locations (Fig. 6). The new scheme reduces the amplitude of the diurnal temperature cycle and produces better night time temperature forecasts. The impact during this season is due to the revised vertical diffusion and the increased skin layer conductivity.

Conclusions

The model changes described above will bring a considerable improvement of the 2m temperature forecasts over land in winter and of the night-time temperatures in summer. They were introduced into operation with the migration of the data assimilation/ forecasting system to the new computer (September 1996). The impact on forecast performance of the large scale weather patterns (e.g. expressed by anomaly correlation of 500 and 1000 hPa fields) is small.

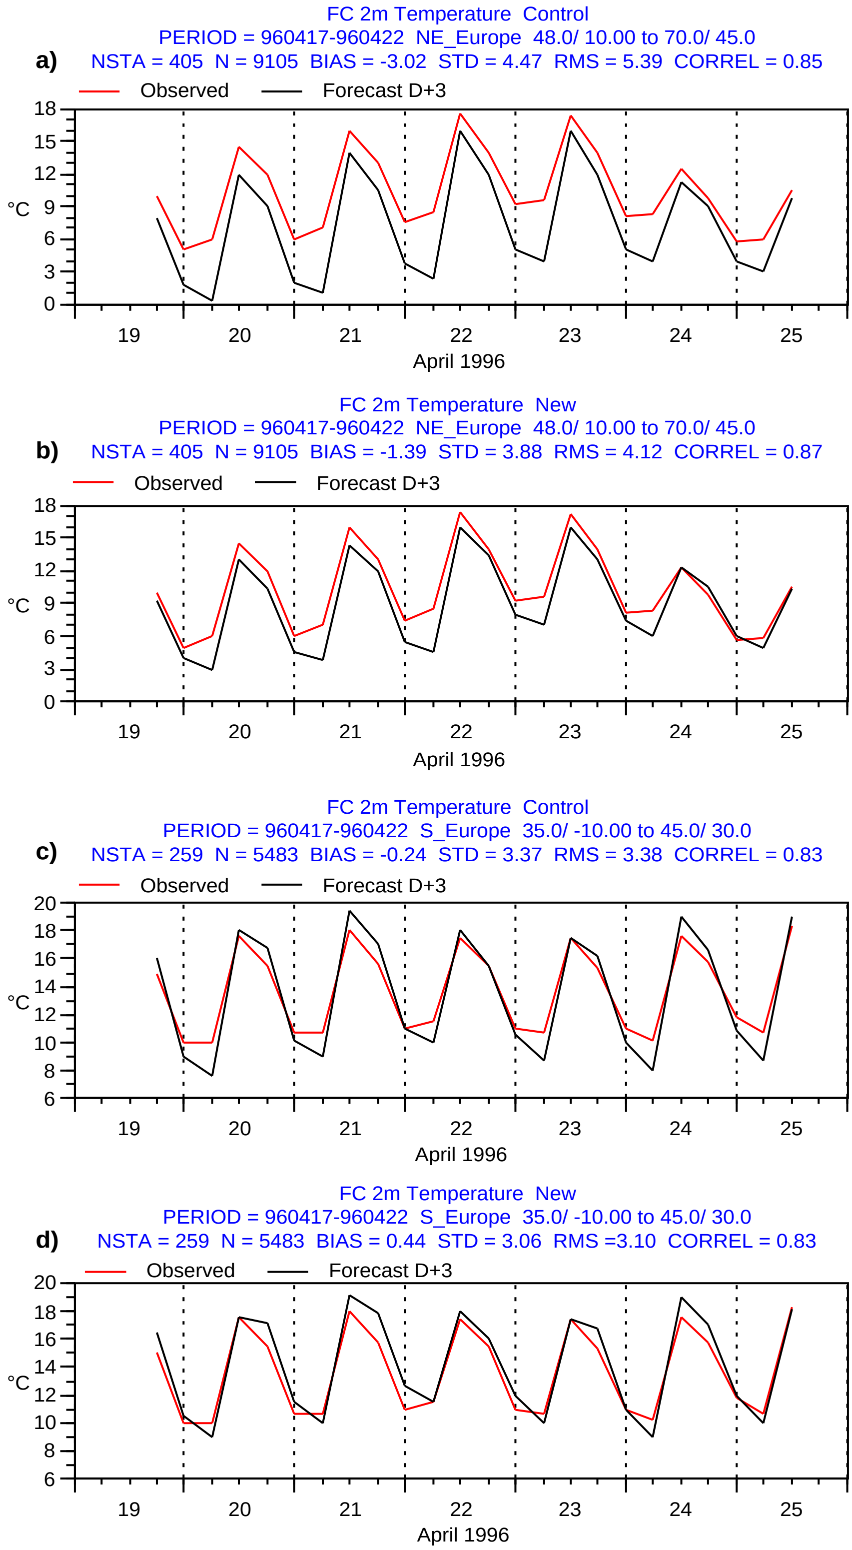

Fig. 6: *Diurnal cycles of 54, 60, 66 and 72 hour forecasts (black) compared with observations (red) for North-East Europe (a,b) and Southern Europe (c,d). These are results from T106 data assimilation/forecast experiments that were initialised with soil temperatures from their corresponding T63 relaxation runs. Figs. a and c correspond to the control experiment, Figs. (b) and (d) correspond to the new model version with revised vertical diffusion, soil moisture freezing and increased skin layer conductivity.*

Anton Beljaars, Jean-François Mahfouf, João Teixeira and Pedro Viterbo

The ECMWF Re-Analysis (ERA) Project

Introduction

The ECMWF Re-Analysis (ERA) project objective was to produce a new, validated 15 year data set of assimilated data for the period 1979 to 1993. During the years of ECMWF's Operational activities the archive of products has been used extensively by researchers all over the world. This archive, in recent years, has supported the ECMWF TOGA (Tropical Ocean Global Atmosphere) data set, providing the global atmospheric data used in much of TOGA research.

The Operational analyses, while providing a valuable resource for such work, is affected by the major changes in models, analysis technique, assimilation, and data usage which are essential to the evolution of continuously improving forecast systems. They also can make use only of those data which become available within near real time. Bengtsson and Shukla expressed the idea that such considerations provide valid reasons for performing a consistent re-analysis of atmospheric data as early as 1988. Typical research applications which could make good use of re-analyses include general circulation diagnostics, atmospheric low-frequency variability, the global hydrological and energy cycle, studies of predictability, coupled ocean-atmosphere modelling, and observing system performance.

ERA only became a practical reality through the funding and assistance received from many quarters, including the ECMWF Council, the European Union, the University of California (PCMDI), the Japan Meteorological Agency (JMA), the World Climate Research Programme (WCRP) of the World Meteorological Organisation (WMO), the Centre for Ocean-Land-Atmosphere Studies (COLA), the National Center for Atmospheric Research (NCAR), the National Centers for Environmental Prediction (NCEP) and Cray Research Incorporated.

David Burridge is the Project Chairman. The project team is Rex Gibson (Project Manager), Per Kållberg, and Sakari Uppala. This team was augmented by Atsushi Nomura from JMA (March 1993 to March 1995), and by two European Union Fellowship Students, Angeles Hernandez (November 1993 to December 1995), and Encarna Serrano Mendoza (January 1994 to April 1996).

In the planning and development phase a Steering Group advised on matters of scientific and policy importance, and additional advice was obtained by setting up an External Advisory Group, comprised of eminent scientists from Europe and the United States of America.

The project began in February 1993. The first phase of the work required the acquisition and preparation of the observations and forcing fields, experimentation to determine the composition of the production system, and the development of both the production system and the internal validation tools. The final production system was adopted in 1994, and there followed a period of sustained production, monitoring and validation throughout 1995 and the first nine months of 1996.

Selection of a "best possible" assimilation system for re-analysis

Before commencing the actual ERA production, it was necessary to define the assimilation system. The aim was to have a proven, but modern data assimilation, not necessarily identical to the then operational suite. The project work thus began in 1993 with a comprehensive set of experiments in the form of parallel data assimilations and forecasts, usually over three week periods and with extensive diagnostics.

The only predetermined limitation was that the re-analyses should be carried out with a horizontal resolution of T106, this was for reasons of resources. For the vertical resolution a test assimilation made us decide upon 31 levels, rather than 19, since the higher resolution produced clearly superior analyses particularly around the tropopause.

A comprehensive test of the "First Guess at Appropriate Time" (FGAT) algorithm was also carried out since FGAT was considered quite expensive in computing time and caused a rather cumbersome file handling. FGAT was indeed shown to give a marginally positive impact in the test, but the benefit was not considered large enough to justify the increased costs.

At the time the experimentation took place, envelope orography was being used in daily operations to parametrize the effects of sub-grid scale mountains. A new parametrization of the effects of sub-grid scale orography based on mean orography, and including a revised formulation of the gravity wave drag, developed by Lott and Miller, was also available. Test assimilations using this scheme showed no negative effects, while up to 10-15% more observations were accepted at 1000 hPa and 925 hPa. In consequence, this scheme was chosen.

Using a prescribed soil climatology, as in the pre-1995 operational system, has the obvious risk of forcing a re-analysis towards its climate, which is based on very sparse information and may suffer from inconsistencies. Hence the new four level self-contained soil parametrization scheme developed by Beljaars & Viterbo for operational implementation was also selected for ERA.

Ongoing work in the Research Department on a new variational assimilation scheme (3D-Var), and a new cloud parametrization with cloud water and cloud fraction as predictive parameters, were not sufficiently mature at the time of decision and were not selected for the re-analysis. The ERA production then began in earnest in December 1994.

However ...

The spring time warm surface temperature biases found in the April 1994 ECMWF operations, were also noticed in the re-analyses of the first year, 1979. These biases were attributed to deficiencies in the cloud cover, particularly over land in spring time, leading to unrealistic drying and heating of the soil. As a quick remedy a very simple soil moisture nudging scheme was introduced, in which the near-ground atmospheric humidity analyses are used to adjust the soil moisture. Production was stopped and new assimilation experiments were run in cooperation with the Research Department in order to test and tune the scheme. Although the nudging reduced the bias problem, the more fundamental problem of too few clouds was still apparent. Since in the meantime the new prognostic cloud scheme of Tiedtke and Jakob had approached maturity and had been shown to yield more realistic clouds, it was decided to run parallel assimilations for July 1985 and January 1986 with the new cloud scheme. The results were so encouraging that it was decided to go back and restart the entire re-analysis from the beginning of 1979 with the prognostic clouds included.

The two complete 1979 assimilations, with the old and the new cloud schemes are quite different, not only in the actual cloud amounts, but also in important derived quantities such as surface energy fluxes and precipitation. This one-year sensitivity 'experiment' demonstrates the very important fact that many aspects of a re-analysis are to a very high degree defined by the assimilating forecast model, and particularly by its physical parametrization.

The ECMWF Re-analysis Data Assimilation System

The principal source of observations has been the ECMWF real-time data collection from the Meteorological Archive and Retrieval System (MARS). Additional sources include:

- 250 km cloud cleared radiance (CCR) data, augmented by NESDIS I-b data where there are gaps;
- Ship and buoy observations from the Comprehensive Ocean Atmosphere Data Set (COADS);
- First GARP Global Experiment (FGGE) and Alpine Experiment (ALPEX) II-b data;
- GMS satellite cloud winds was made available by the JMA;
- PAOBS data from NMC Melbourne;
- The JMA also provided supplementary radiosonde and aircraft data;
- Subduction and TOGA buoy data.

The ERA assimilation system is based on:

- Spectral T106 resolution with 31 vertical hybrid levels;
- Intermittent statistical (Optimum Interpolation) analysis with 6 hour cycling and no FGAT;
- One dimensional variational (1D-Var) physical retrieval of TOVS cloud cleared radiances;
- Diabatic, non-linear normal mode initialisation (5 vertical modes);
- ECMWF's Integrated Forecast System (IFS) forecast model with a fully 3D semi-Lagrangian advection scheme;
- Mean orography with a compatible parametrization of the effects of sub-grid scale orography;
- A four-layer prognostic soil scheme, with "nudging" of the soil moisture from boundary layer atmospheric humidity;
- ECMWF mass-flux convection;
- Prognostic equations for cloud water content and cloud cover;
- ECMWF operational planetary boundary layer parametrization;
- ECMWF operational radiation parametrization with prescribed concentrations of aerosols, carbon dioxide and ozone, with the ozone varying geographically and seasonally, the aerosols varying geographically and vertically, and carbon dioxide held constant.

The only other externally prescribed forcing comes from the sea surface temperatures (SST) and sea ice cover. These were obtained from:

- For December 1978, and from January 1980 to October 1981 the UK Meteorological Office (GISST) 1 degree monthly SST analyses;
- For the final rerun of 1979 the EOF-enhanced SST from the combined UK Meteorological Office (GISST) and NCEP effort;
- NCEP (Reynolds) 1 degree weekly OI SST analyses (November 1981 to 1993);
- Ice limits derived by A. Nomura of the ERA group from SMMR and SSM/I data.

Special arrangements have been made with NCEP and DAO to enable each of the three re-analyses to use an identical set of observations, sea surface temperatures, and ice limits for 1979.

Apart from the resolution, the ERA assimilation system is identical to that used for ECMWF operations from April 1995 until February 1996.

ERA Production

It was essential to develop a reliable production system capable of performing data assimilation very quickly. Using the combined experience of the Centre's Operations and Research Departments, the systems in use were studied carefully, slimmed down where necessary, modified to exploit a data archive as opposed to real-time data, and optimised for performance. This resulted in a prototype system capable of performing at the required rate, which was further refined and completed while being used as the principle vehicle for the initial ERA experimentation.

Once production began, late in 1994, the scientific emphasis gradually moved from experimentation to monitoring and validation. The external forcing fields were validated before the production started by means of maps, averages and time series. Every effort was made to detect, as early as possible, potential problems which would require further investigation. Where appropriate production was halted and restarted from an earlier date; in some cases production was allowed to continue, but the month or months concerned rerun at a later time. The production monitoring made use of a set of quality control tools, whose output, usually in the form of graphical information, was monitored continuously. All graphical and tabular monitoring results were kept both as hard copies and as files. Diaries were kept of all special events and problems encountered. Some of the monitoring performed is described in subsequent sections.

Production and monitoring continue throughout 1995 and into 1996. During the second quarter of 1996 the first pass through the full 15 years was completed. Monitoring enabled many errors which happened during the production to be located and rectified; nevertheless two lengthy periods needed to be rerun. First, the early production (to August 1980) was adversely affected by a bug, present also in the operational system, which significantly affected humidity at upper levels. Secondly, all cloud track wind data (SATOB) except the GMS data were accidentally excluded from June 1990 to October 1992, due to a change in their format within the ECMWF archive. A rerun of the first period was particularly desirable, as it presented an opportunity to run the FGGE year with the same observations and forcing fields as NCEP. Both reruns were completed in September 1996.

The internal validation programme was augmented by setting up a number of external validation projects with the following partners:

- Max Plank Institut für Meteorologie, Hamburg, Germany (MPI)
- Hadley Centre, Meteorological Office, Bracknell, U.K. (UKMO)
- Météo-France, Toulouse, France (DMN)
- Danish Meteorological Institute, Copenhagen, Denmark (DMI)
- Royal Netherlands Meteorological Institute, De Bilt, Netherlands (KNMI)
- Instituto per lo studio delle Metodologie Geofisiche Ambientali, CNR, Modena, Italy (IMGA)
- Atmospheric Dynamics Group, University of Bologna, Bologna, Italy (ADGB)

The continuous exchange of information between the ERA team and the validation partners has proved extremely useful. It will result in material being available to potential future users of the data which should be of considerable assistance.

Close cooperation has also been established between the ERA team and the teams responsible for the NCEP re-analysis, and the re-analysis performed by the NASA Data Assimilation Office.

Observation Monitoring

Within each month the characteristics of observations were fixed with few exceptions. The use of observations, and the exclusion of undesirable data were controlled by means of "blacklists". For each month of production there exists a corresponding blacklist file. From the observation diagnostics alone (or combined with analysis monitoring results) undesirable observations were manually blacklisted for subsequent months. The following is a list of the main observation monitoring diagnostics routinely produced during ERA:

- Total counts of all observation types for each analysis cycle.
- Observation distribution maps for each analysis cycle.
- Reject observation distribution maps for each analysis cycle.
- Vertical profiles of observation fit to first guess, uninitialised analysis and initialised analysis for each analysis cycle.
- Monthly time-series graphs of CCR brightness temperatures for all channels.
- Monthly time-series graphs of CCR brightness temperature - first guess for all channels.
- Once monthly maps of geographical distribution of CCR bias correction fields.
- Monthly maps of radiosonde - first guess biases.
- Vertical profiles of monthly mean and standard deviation of radiosonde biases, for bias-corrected areas.
- Monthly maps of mean & rms errors of drifting buoys

Radiosonde Height Biases

In ECMWF operations a comprehensive system for the correction of systematic errors (biases) in radiosonde temperature and geopotential is used. It is based on feedback information from the data assimilation and determines corrections, which are a function of the solar elevation, to individual radiosonde observations made by known instrument types. A simplified version of this scheme has been adapted for ERA. Bias corrections are calculated from the preceding twelve months of observation minus first-guess statistics for those WMO blocks where homogeneous radiosonde equipment were believed to have been used. The correction is then applied above 250 hPa.

Continuous manual monitoring of analysis feedback information and monthly mean analysis increment fields reveals areas (stations) with large and suspicious increments. Suspicious areas are added to the bias correction scheme, and individual stations with very large first-guess deviations are added to the blacklist. The monitoring is also applied to bias-corrected and blacklisted stations so they can be removed from the correction when they improve. There is thus a continuously developing radiosonde bias correction and blacklisting. All information on correction and blacklisting is kept.

The correction has been applied through the whole ERA period, however, the corrections calculated from 1980 were used at the start of the 1979 rerun. This assumes that instrument characteristics did not significantly change between these years. The correction is largest in the beginning of the period and decreases clearly towards 1989 and from 1990 only small corrections are applied to a small number of areas.

TOVS Quality Control and Bias Tuning

The CCRs constitute a consistent data set for the entire re-analysis period. Before 1D-Var retrievals can be calculated the observed radiances have to be tuned with the forecast model. This tuning corrects systematic errors in the forecast model and possible errors of the forward model. Monthly corrections are determined for each channel by comparisons with first-guess forecasts in the vicinity of a selection of quality controlled radiosonde

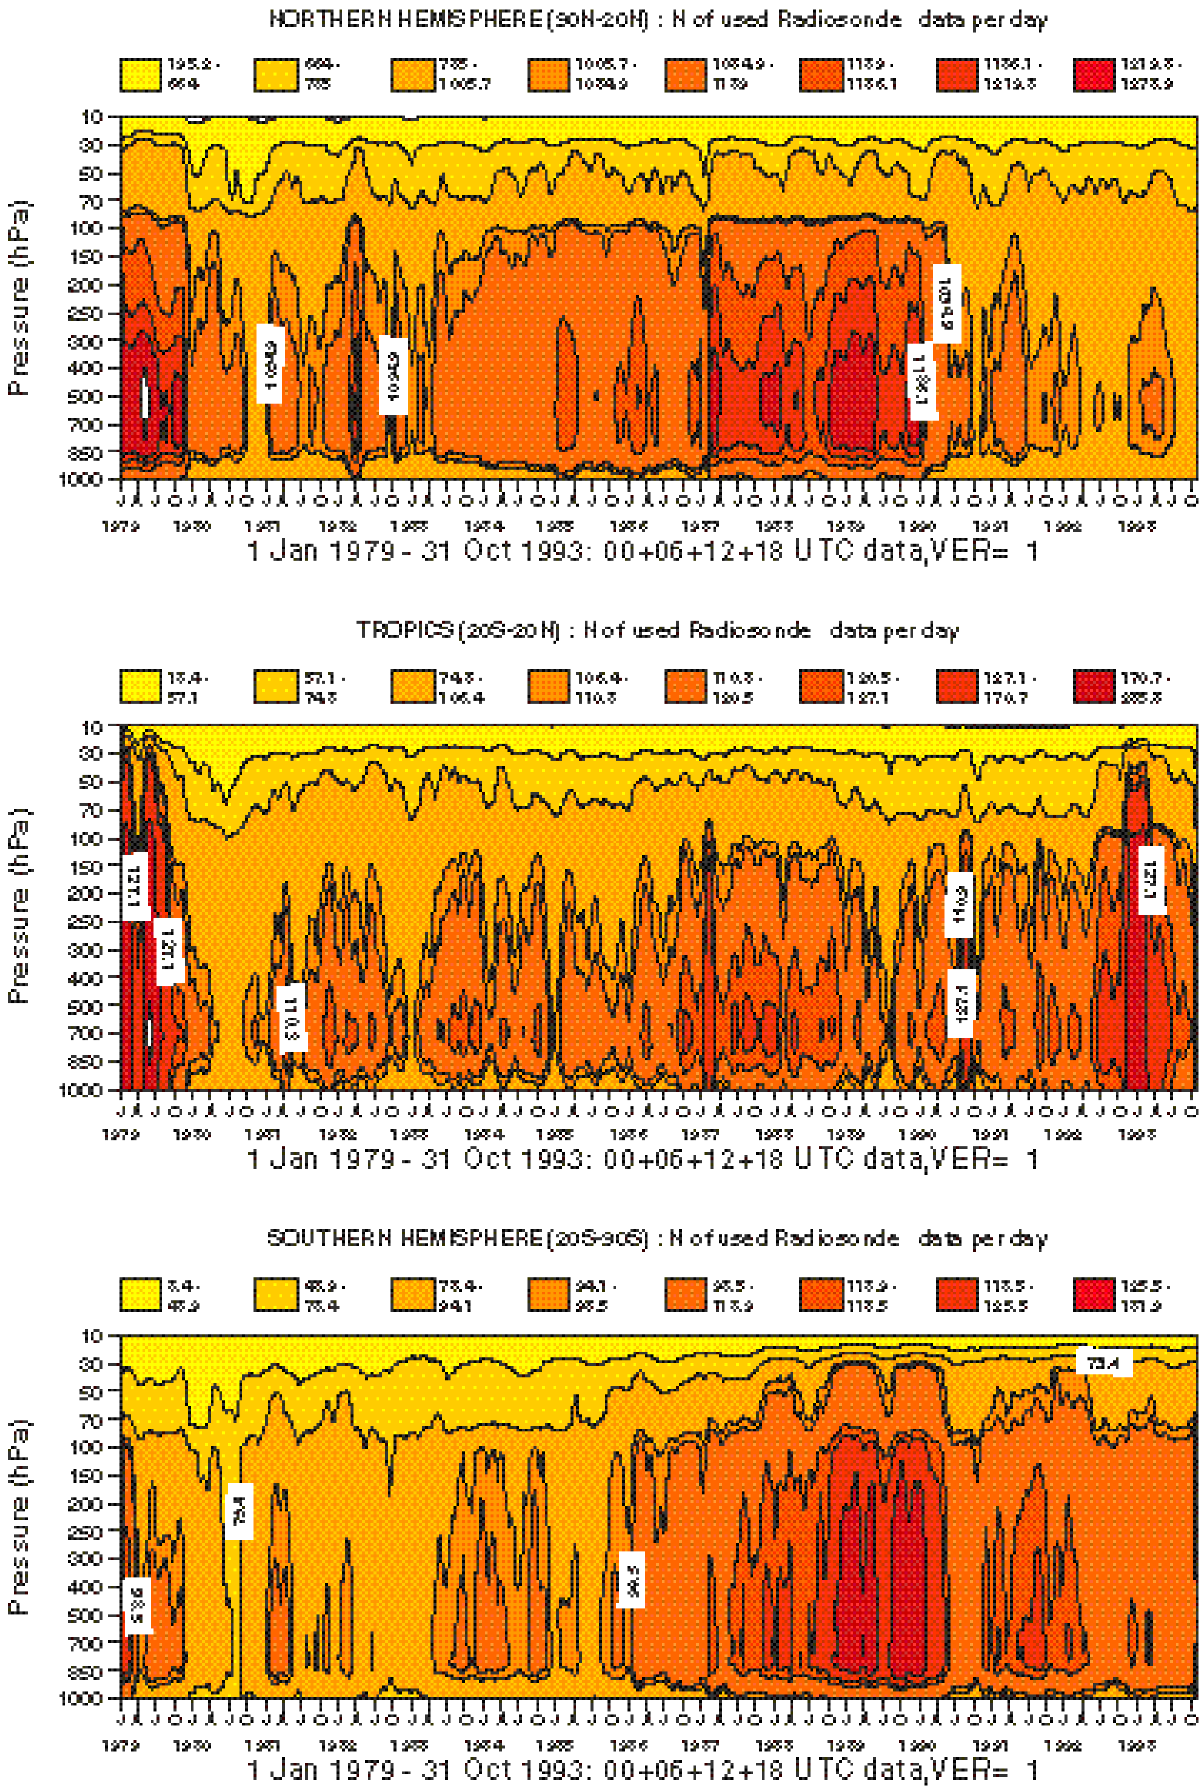

Fig. 1: *The variability of the number of radiosonde data used daily in the reanalysis over the period 1979 to 1993.*

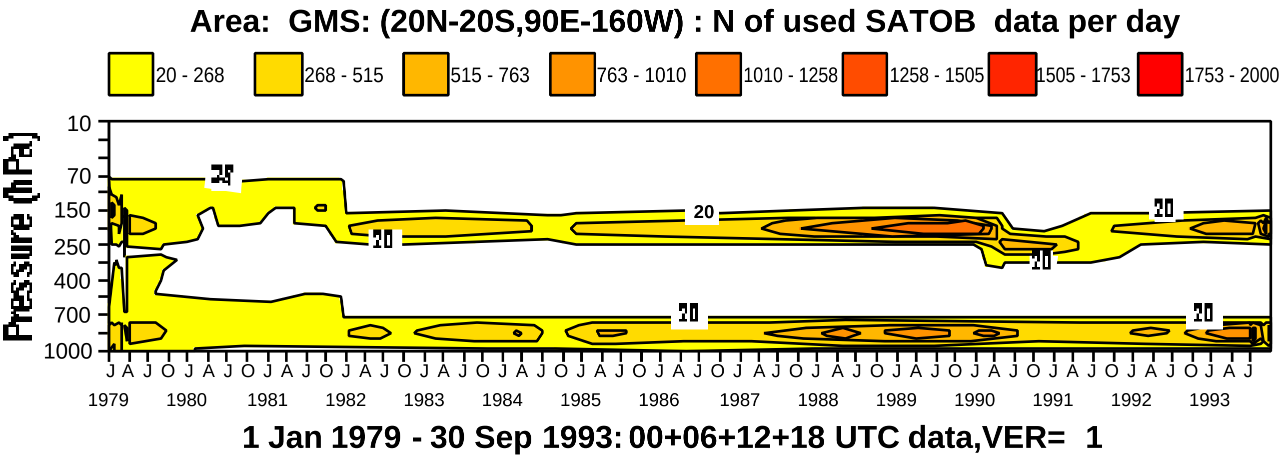

Fig. 2: *The variability of the number of GMS cloud drift wind data used daily in the reanalysis over the period 1979 to 1993.*

data in different parts of the world. These corrections are then applied during the following month of assimilation.

The first step of the quality control is carried out before the data are used in the analysis. For each satellite a monthly time series of the mean six hour brightness temperatures for each channel are plotted together with the corresponding number of data. This is important for several reasons. It is good to know beforehand if and when a channel drops out and when it comes back, both for the automatic bias tuning, and also for the calculation of a set of 'physical' 1D-Var retrievals. Information obtained from NOAA/NESDIS - Polar Orbiter Archived TOVS Sounding Data change and Problem Record - contains a list of the most important changes or errors (e.g. miss-locations of data) as well as changes in the software used to calculate the CCR data. Most of these events are difficult to identify from brightness temperatures alone, since their effects are relatively small.

The derivation of 1D-Var is performed using the bias corrected brightness temperatures. For quality control purposes the mean corrected and uncorrected observation minus first guess departures for each six hour period are plotted as a monthly "radgram" for each satellite and channel. Since the first guess itself is independent of any CCR change, at least when a change is about to happen, these graphs reveal satellite problems that have occurred during the previous month. Often NESDIS has listed a change (e.g. a change in the water vapour attenuation coefficients), but it is only afterwards that it can be seen whether or not this change has caused a significant problem in the data assimilation. In practice full use of this information would require the bias tuning to be done separately during all the abnormal periods, and those periods subsequently rerun with the new coefficients.

Monitoring of Analysis Quality

Data-assimilation works well if the analysis increments (analysis minus first guess) are small compared with the changes that the atmosphere experiences on average from one cycle to the next. In addition to a set of 15 synoptic analysis maps, which were produced at each analysis cycle, the following diagnostics were produced:

- A set of 13 synoptic analysis increment maps at each analysis cycle.
- Monthly time series plots of the global mean and standard deviation of the cycle analysis increments.
- Monthly maps of the temporal mean and standard deviation of the analysis increments.

Observing System Characteristics

During the re-analysis period the global observing system experienced many changes. There were periods with one polar orbiting satellite, and others with two. Figure 1 shows the variability of the number of radiosonde height data, which the analysis has accepted and used. The FGGE year stands out clearly, as does the period 1987 to 1990. The year 1980 was a low point; from then on data slowly increased until 1986. After 1990 the radiosonde network slowly deteriorated expect in the tropics. The way observations are derived has also changed. Figure 2 shows as an example the time evolution of the number of used GMS winds. The different height assignment algorithms have produced winds on different altitudes. The number of aircraft winds also has increased through the years, and the number of drifting buoys has varied considerably.

Figure 3 shows the time evolution of the root mean square differences between the first-guess and the radiosonde heights through the re-analysis period in the Northern Hemisphere. It was produced using the early data for 1979 which did not include the radiosonde bias corrections. The improvement from 1979 to 1980 is significant above 250 hPa and is mostly due to the bias correction of radiosonde heights. There are still signs of steady improvement from 1980 onwards, which especially towards the end of the period is due to the improved quality of radiosondes.

The TOVS CCR data set is relatively uniform throughout the re-analysis period; the same instruments, HIRS, MSU and SSU, were used with the same specifications on all satellites. The NESDIS temperature and humidity retrievals used historically in ECMWF operations vary

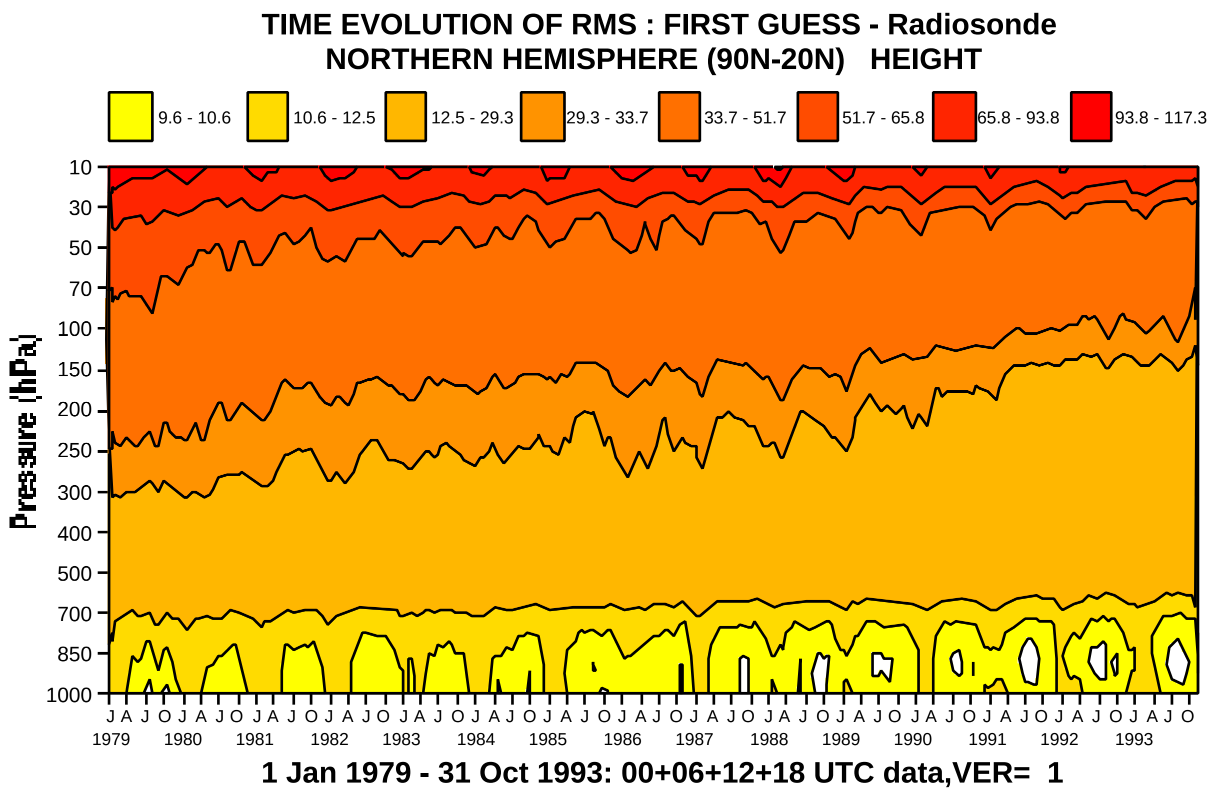

Fig. 3: The evolution of the RMS of the difference between first-guess heights and radiosonde heights in the Northern Hemisphere for the period 1979 to 1993: Units m.

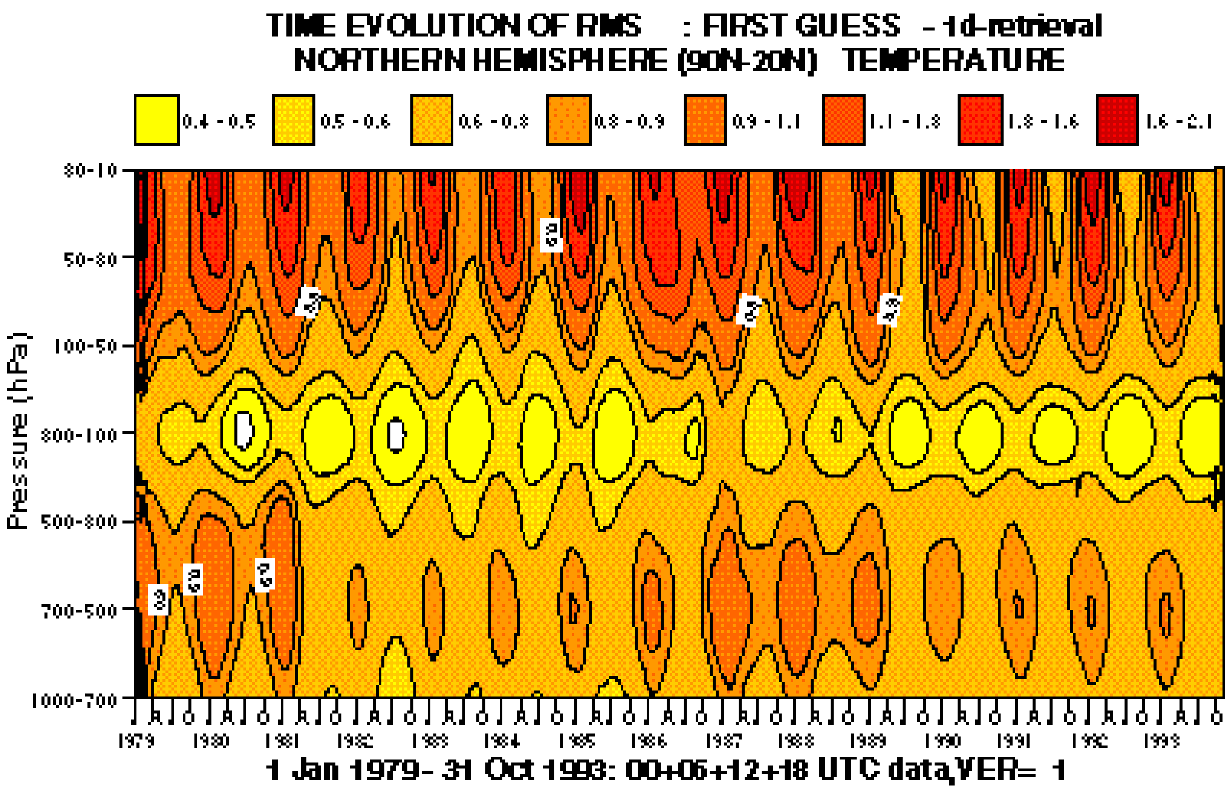

Fig. 4: The evolution of the RMS of the difference between first-guess layer mean temperatures and 1D-Var retrieved layer mean temperatures in the Northern Hemisphere for the period 1979 to 1993: Unit °C.

considerably with time due to ongoing development and tuning of the statistical retrieval algorithms. As can be seen from the time-evolution of the mean RMS of 1D-Var layer mean temperatures in the Northern Hemisphere, Figure 4, the two periods with one satellite, January 1979 to August 1981 and mid 1985 until the end of 1988, shows slightly larger errors than the rest of period with two satellites.

The quality of cloud drift winds has improved also as can be seen from the time evolution of the wind RMS of first guess minus observed values in Figure 5. The RMS values of 9 m/s in 1979 has been reduced to 6 m/s in 1993.

There is overall evidence that the global observing system improves through the period.

Some general circulation features in ERA Comparisons with ECMWF operational analyses

Tropical divergence

The operational analysis and forecasting system at ECMWF has undergone continuous development and upgrading through the years, this is indeed the very reason for doing a re-analysis at all. During our monitoring and validation we have compared different aspects in ERA with the old operations and we will show an example from 1983, when the operational model was an N48 grid point model with a resolution of 1.875° by 1.875° and with 16 vertical levels. The parametrization at that time included a diagnostic cloud scheme for the radiation and a Kuo-type convection. The SST analyses used were the operational NMC analyses received at ECMWF in a 5° by 5° resolution. This was near the end of the ECMWF 'grid point era', the spectral T63 model was introduced in April 1983. Diabatic initialization had been introduced in September 1982. Furthermore, in 1983 the TOVS were used in the form of NESDIS SATEM with a resolution of 500 km rather than the 250 cloud cleared radiances used in ERA.

The two maps in Figures 6 and 7 show the monthly mean velocity potential and divergent wind at 150 hPa during March 1983. This was during the most intense phase of the 1982-83 El Niño/ Southern Oscillation (ENSO) event. The anomalous sea surface temperatures, up to 4-5 °C warmer at the equator around 140°W, forces the major convective area to migrate from its normal position east of New Guinea to the position in Figure 6. As vividly seen in Figure 7, the ECMWF operational analyses at the time were unable to catch this anomalous circulation, indeed there is very little divergent outflow to be seen at all.

The 1982-83 ENSO event shows up in many other aspects of the ERA circulation as well. The famous pressure oscillation between Papeete (Tahiti) and Darwin (northern Australia) is well captured, as are the draughts in tropical Africa and northeastern Brazil and the excessive precipitation over the central Pacific. At higher latitudes the Pacific-North American anomaly (PNA) pattern in the surface pressure which is a fingerprint of anomalous circulation over large parts of the Northern Hemisphere is clearly seen. Some readers may remember that the winter of 1983 was warm in northeastern Europe and cold in the southwestern parts. The ERA mean maps show how this anomaly was connected with the ENSO event through the PNA.

Tropical cyclones

All tropical cyclones in ERA were tracked, using vorticity maxima at 850 hPa and surface pressure minima in the relevant areas and seasons. The tracks were compared with existing 'best track' data based on satellite imagery and airborne in situ observations. Over the 15 years ERA is able to find about 70% of all reported tropical disturbances and cyclones, albeit not at all with the observed intensity. The RMS positional error is of the order 150 km, not much more than one grid point. The ECMWF operations reached a similar accuracy only towards the end of the 1980's.

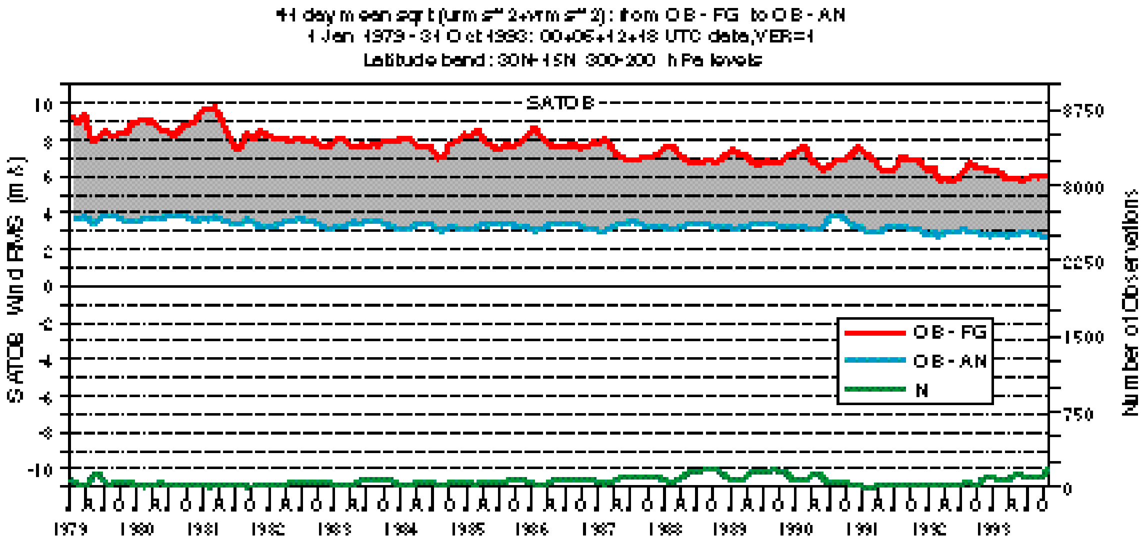

Fig. 5: *The evolution of the RMS of the difference between first-guess wind vectors and upper level cloud drift wind vectors (red line), and the difference between analyzed wind vectors and upper level cloud drift wind vectors (blue line) in the band (30° N-15° N for the period 1979 to 1993: Units m s*$^{-1}$*. The green line is the number of data used over the period.*

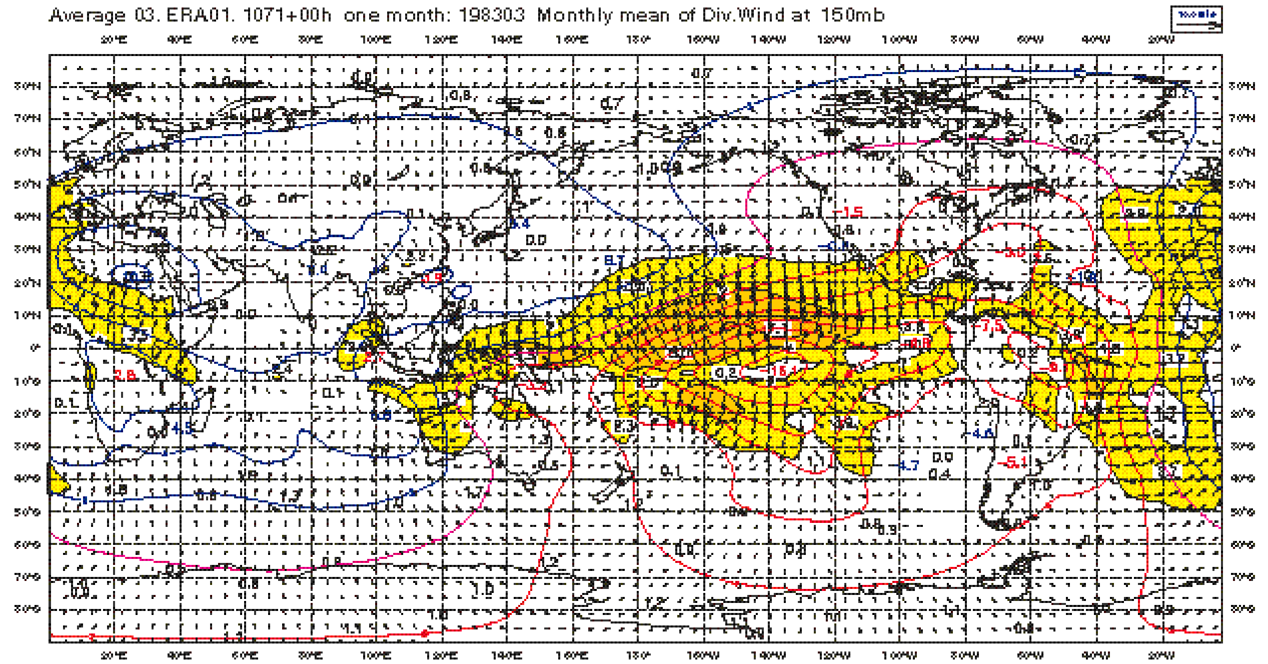

***Fig. 6:** The monthly mean velocity potential at 150 hPa (blue and red contours) from the reanalyses for March 1983: Units 10^{-5} m^2 s^{-1} The divergent wind at this level is illustrated by the arrows.*

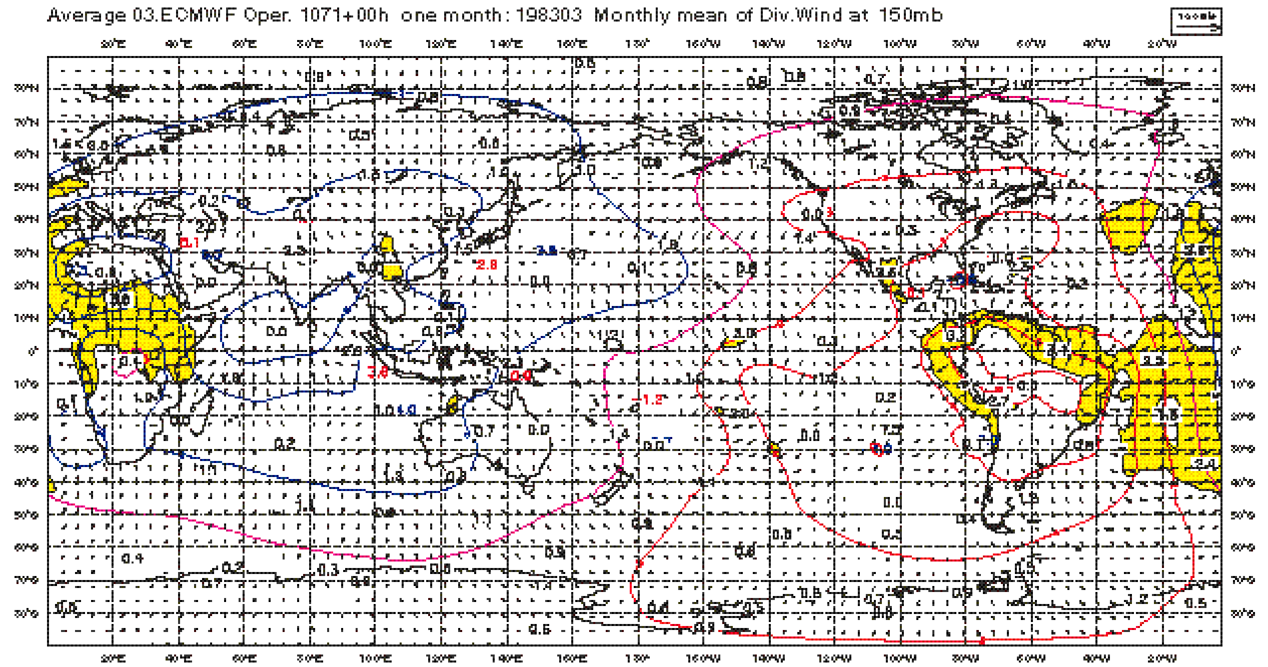

***Fig. 7:** The monthly mean velocity potential at 150 hPa (blue and red contours) from the operational analyses for March 1983: Units 10^{-5} m^2 s^{-1}. The divergent wind at this level is illustrated by the arrows.*

Looking at the performance in the different ocean basins, the picture is somewhat more mixed. In the eastern Pacific the tropical cyclones are poorly analysed both in number and in location, this is an area with very few observations. In the Atlantic and the western Pacific, which has the greatest number of tropical cyclones, they are analysed quite well. In the northern Indian Ocean and the Southern Hemisphere it is difficult to draw statistically safe conclusions due to too few tropical cyclones.

The stratosphere

The ERA assimilation system has only four levels above 100 hPa, at 70, 50, 30 and 10 hPa respectively. The potential of the analyses to capture stratospheric circulation features such as the quasi-biennial oscillation (QBO) was thus not assured. To make things worse, near the equator there are very few reliable radiosondes reaching 10 hPa or above, and the NESDIS TOVS data used elsewhere above 100 hPa in ERA were not used between

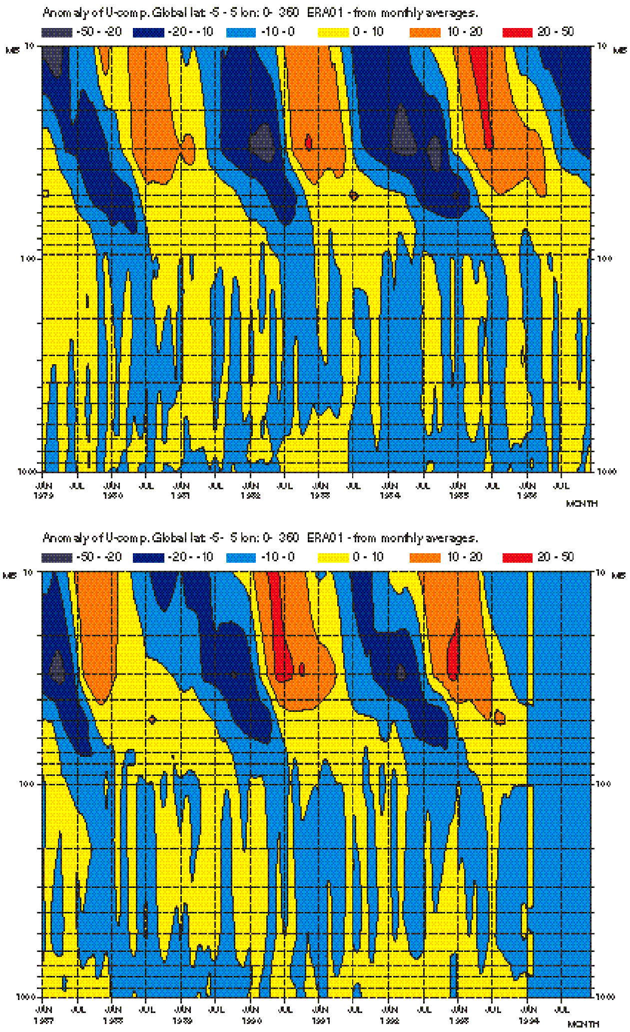

Fig. 8: The evolution of the monthly anomaly of the mean zonal wind between 5° N and 5° S over the period 1979 to 1993: Units m s^{-1}.

20°N and 20°S just because of their bad vertical resolution which, in tests, was found to smear out the strong vertical wind shear of the QBO phenomenon.

In spite of all these 'ifs and buts', the re-analyses do exhibit a very well developed QBO. In Figure 8 a time-height diagram of the monthly mean zonal wind anomalies is shown. The quasi-biennial switch between easterlies and westerlies, and the downward migration of the anomalies is clearly seen. Similar plots localised over Singapore (not shown here) confirm that the ERA QBO is quite realistic compared to the soundings.

Another stratospheric feature that shows up beautifully in the re-analyses is the global heating in the stratosphere caused by the aerosol clouds emitted by the volcanoes El Chichon (Mexico) in 1982 and Pinatubo (Philippines) in 1991. Since the assimilating model does not know about these emissions, the global temperature anomalies seen in Figure 9 is entirely due to the data, in this case primarily the TOVS temperatures and some radiosondes. In the figure one can also notice a gradual cooling of the stratosphere over the 15 years.

Surface fluxes

Good estimates of the surface fluxes of energy, momentum and water, over land and particularly over the oceans, are of fundamental importance for the understanding and modelling of the coupled atmosphere-ocean system. Estimates of these fluxes from observations is difficult and existing attempts vary considerably in coverage and quality. Thus the user community has expressed great expectations from the re-analyses, where an assimilating model actually calculates them globally and in detail. In ERA the fluxes are extracted from twice daily forecasts up to +24 hours.

It is well known that model generated fluxes generally suffer from spin-up or spin-down, i.e. they increase (or decrease) with the forecast length. This is also true for the ERA fluxes. Figure 10 shows the zonally averaged total precipitation for four forecast lengths, 00->06 (cyan), 00->12 (green), 00->24 (red) and 12->24 (blue) hours. It is evident that the precipitation intensity increases with the forecast length during the first 24 hours in these 11-year averages. Other fluxes, such as the evaporation and the net energy also exhibit spin up problems. Thus, in the 11 year average, the global net energy exchange drifts from 7 W/m^2 going from the atmosphere to the oceans in the 00->06 forecasts to 3 W/m^2 going from the oceans to the atmosphere in the 12->24 forecasts.

There are also very large inter annual variations in the fluxes, some of which may be related to varying data coverage, for instance the loss of satellite wind data over the Indian Ocean in 1980. Others are likely to be indications of real inter-annual variations in the global circulation. The distinction between data coverage and data quality related variations and variations due to genuine circulation changes will be a major, and difficult task for future research based on ERA. Some assistance may be given by the upper air general circulation statistics, such as variances and covariances, that have been collected during the assimilation, but so far not evaluated.

ERA forecasts

Preliminary evaluation of the skill of the ERA forecasts indicates a relatively steady performance through the years and a clear improvement to the corresponding old operational forecasts; the skill equals that of operations in 1993 (Figure 11). Locally, for example over North America or Europe, the skill has variability, which possibly can be related to changes in climate or observing system, but this has not been studied yet.

Data for Research

The analysis and forecast results from ERA production and experimentation are stored, and may be retrieved by internal ECMWF and Member States users using MARS. The Data Services unit provides data for external research scientists. They have already become an important data source for many applications, both internal and external, including most of the potential applications listed in the introduction above.

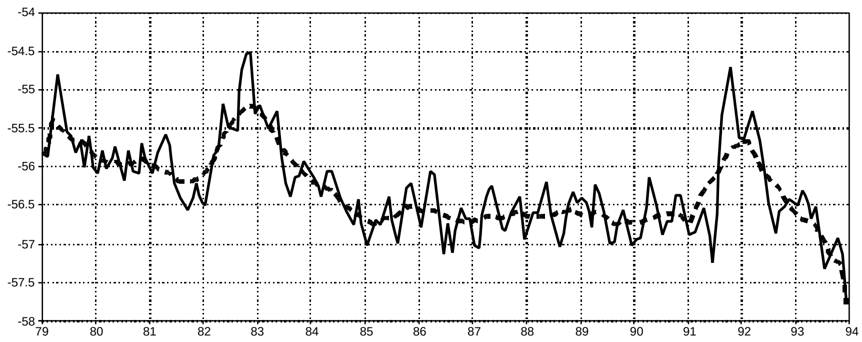

Fig. 9: *The evolution of the global monthly temperature anomaly (full line) and 12 month moving average (dashed line) at 30 hPa for the period 1979 to 1993: Units °C.*

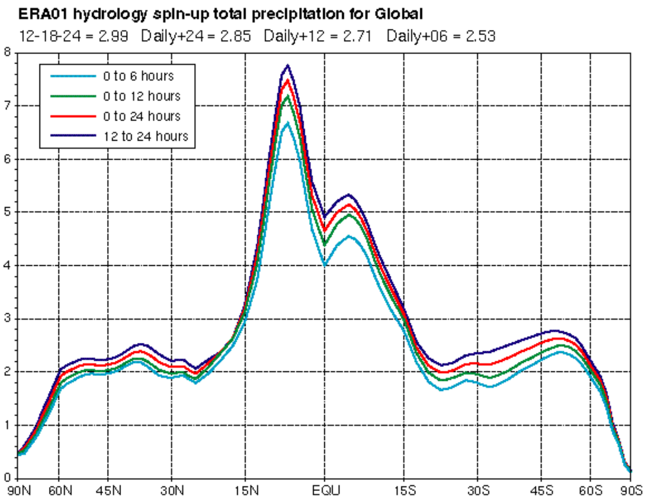

Fig. 10: *Estimates of the zonally averaged total precipitation rate for four forecast periods 0 to 6 hours (cyan), 0 to 12 hours (green), 0 to 24 hours (red) and 12 to 24 hours (blue): Units mm day*$^{-1}$*.*

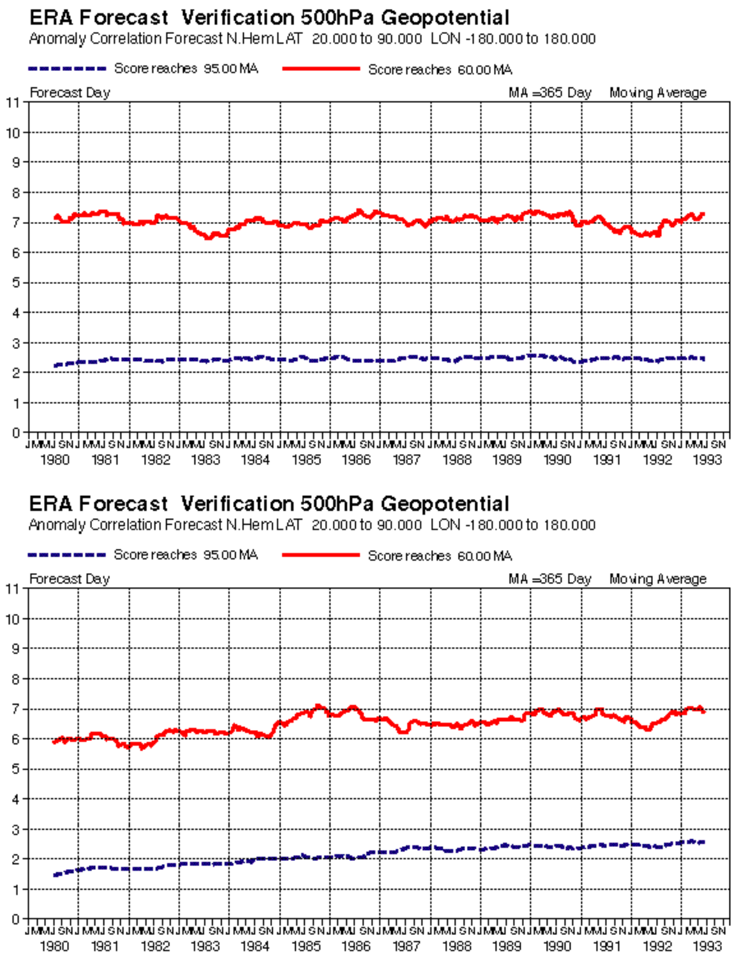

Fig. 11: *The day when the forecast skill (anomaly correlation) drops to 95% (blue) and 60% (red) for the period 1979 to 1993 (a 365 moving average has been applied).*

While the work above gives an introduction to the research made possible through re-analysis, it is really only a beginning. Indeed, there has not been sufficient space here to describe the many additional interesting results emerging both from the Validation Partners, and from other work done within ECMWF.

The ECMWF Data Services provide data from specific data sets within the archive, performing extraction and interpolation where needed, and supplying appropriate software with each order to assist users in decoding the data. Data are supplied in conformance with the ECMWF Council's rules for the provision of data, and charges are made to cover extraction costs. Currently data can be supplied on 1600 or 6250 bpi 9 track tape, IBM 3480 cartridge tapes, or on Exabyte. Supply of data via CD-ROM is under development. A brochure describing the services supported for ERA data, and including a set of order forms, is available on request.

Rex Gibson, Per KållbergSakari Uppala

Extreme Rainfall Prediction using the ECMWF Ensemble Prediction System

Abstract

The combined use of the European Centre for Medium-Range Weather Forecasts (ECMWF) high-resolution at T213 spectral triangular truncation and with 31 vertical levels (T213L31) operational model and Ensemble Prediction System (EPS), during cases of intense Mediterranean storms, is studied. In particular, it is discussed how EPS products can be used to provide a measure of confidence in the high-resolution precipitation forecast. Three case studies (two extreme events plus one false-alarm case) are analyzed. For the first two cases, the EPS probability values for precipitation occurrence supported the medium-range T213L31 prediction, which proved to be successful. By contrast for the third case, the high-resolution forecast suggested heavy rainfall over northern Italy but was not supported by the EPS. The T213L31 prediction for this case was poor. EPS forecasts of extreme weather events are necessarily compromised by the moderate resolution of the T63L19 model (a version of the operational high-resolution T213L31 model, at T63 spectral triangular truncation and with 19 levels) used to generate the ensembles. In future studies, ensembles will be made using at least T106L31 resolution combined with an increase in ensemble size.

1. Introduction

Since December 1992, the European Centre for Medium-Range Weather Forecasts (ECMWF) operational 10-day predictions have included both a high-resolution deterministic forecast and an Ensemble Prediction System (EPS). The development of the ECMWF EPS follows the work of *Epstein* (1969), *Gleeson* (1970), *Fleming* (1971a-b) and *Leith* (1974), who laid the theoretical and numerical foundations of a probabilistic approach to weather forecasting.

At the time of writing, the high-resolution model has a T213 spectral triangular truncation and 31 vertical levels (*Simmons et al*, 1989, *Courtier et al*, 1991), whilst the EPS consists of 33 lower resolution (T63L19) forecasts (*Molteni et al*, 1996), made by perturbing the operational analyses using dynamical singular vectors (*Buizza and Palmer*, 1995).

The purpose of this study is to show how the EPS probability predictions can be used to estimate the degree of confidence that could be associated with a deterministic high-resolution prediction. The discussion focuses particularly on precipitation for three different cases. The study is performed using the type of products available to the ECMWF Member States' Meteorological Operational Services.

The first two cases were characterized by very intense rainfall, which led to floods, loss of lives and enormous damage. For both cases the deterministic prediction of precipitation proved to be very successful and well supported by EPS forecasts of the probability of precipitation exceeding predefined thresholds. By contrast, the EPS did not support the deterministic prediction of heavy rainfall for the third case, for which the operational high-resolution prediction proved to be wrong.

The paper is organized as follows. In Section 2, a brief description of the ECMWF EPS is given, in Section 3, a synoptic overview of the three analyzed cases is reported, while Sections 4, 5 and 6 discuss the deterministic and probabilistic predictions for the three cases. Conclusions are drawn in Section 7.

2. ECMWF Ensemble Prediction System

The ECMWF EPS comprises, at the moment of writing, 32 perturbed and one unperturbed (control) non-linear integrations of a version of the ECMWF model (*Simmons et al*, 1989, and *Courtier et al*, 1991) with spectral truncation T63 and 19 vertical levels. The initial conditions are created by adding and subtracting 16 perturbations to the control initial conditions. The initial perturbations are defined using the singular vectors (*Buizza and Palmer*, 1995) of a linear approximation of the ECMWF model. The singular vectors (SVs) identify the most unstable directions of the system growing over a finite time interval (48 hours) named the optimisation time interval. Once the perturbations have been constructed (as linear combinations of 16 selected singular vectors), they are added and subtracted to the control initial conditions to define 32 perturbed initial conditions. Then, 32 plus 1 (control) 10-day T63L19 non-linear integrations are performed. With the current ECMWF computer facilities (CRAY C90 with 16 processors), the total elapsed time is approximately 2.7 hours, which is about 1.3 times the elapsed time needed to performed the 10-day T213L31 ECMWF operational forecast.

3. Synoptic overview

3.1 The Italian storm: 4-6 November 1994

During the first week of November 1994, heavy rainfall caused catastrophic floods and land slides over southern France and northern Italy. The north-western Italian Piedmont was the region which suffered the most, with the loss of more than 60 lives.

A detailed description of the evolution and spatial variability of the observed precipitation can be found in Buzzi et al (1995), who identified the period between the 12 UTC of the 5th and 6th of November as the most critical (also see Athanassiadou and Thorpe, 1995, Binder and Rossa, 1995, Buzzi and Tartaglione, 1995, and Dorninger and Hantel, 1995). Our analysis will concentrate on this time interval.

Between the 2nd and the 4th of November 1994 an upper-level cyclonic circulation system was propagating across the northern Atlantic with an associated trough deepening east of Iberia. The associated low-level advection of warm and moist Mediterranean air-masses led to moderate and heavy rain-falls over Provence, the French and Italian Alps, and the western Po valley from the 2nd of November.

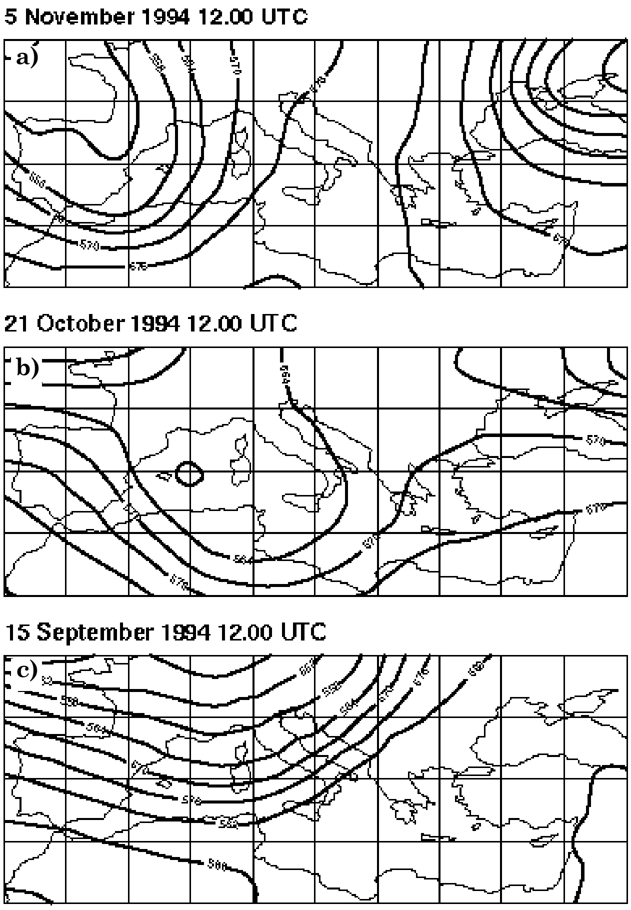

Fig. 1: *a) ECMWF 500 hPa geopotential height analysis for 94.11.05 12.00 UTC. b) as a) but for 94.10.21. c) as a) but for 94.09.15. Contour interval: 60m.*

On the 5th of November 1994, while the eastward progression of the upper-level low slowed and eventually became quasi-stationary off Ireland, the upper-level trough moved across Iberia into the western Mediterranean (Fig. 1a). Meanwhile, low-level advection of warm moist air continued. On the 6th of November, a cut-off formed over the Baleares and moved across Corsica and northern Italy to the eastern Alps. As a result of the progression of cold air into the western Mediterranean, of the sharpening of air-mass gradients and of the formation of the cut-off low, frontal activity and associated rainfalls were substantially enhanced.

3.2 The Hellenic storm: 20-22 October 1994

Between the 21st and the 22nd of October 1994, heavy rainfall caused major disruption in many areas of the Hellenic peninsula and more specifically in the greater area of Athens. The catastrophe caused the loss of twelve lives and significant damage to property (for more details see Lagouvardos et al, 1995).

On the 20th of October 1994, an upper level trough (with a NW-SE tilt) associated with a deep low off Ireland was penetrating south-ward over western Europe, leading to a cut-off low over the western Mediterranean on the 21st of October 1994 (Fig. 1b). On the forward side of this system low-level warm moist air was advected over the Hellenic peninsula. During the following days the low circulation system moved slowly east-wards and the associated continuous convergence over the Greek mainland of air-masses originating from the central and southeastern Mediterranean was enhanced. The combined effect of the convergence of the warm moist air-masses and their interaction with colder air-masses converging from the northeast (being advected on the southern side of a continental high pressure system), resulted in extreme rainfalls over many areas of Greece.

3.3 The false-alarm case: 14-16 September 1994

On 14 September 1994, a deep low was centred over north-western France. On its southern side Mediterranean air-masses were advected over Italy towards central Europe. During the following 24 hours the low circulation system deepened and moved east-wards with its centre over Denmark starting to fill slowly. Westerly, slightly cyclonic flow prevailed over the western Mediterranean area (Fig. 1c).

Considering the 24 hour time interval (between 12 UTC of the 15th and the 16th of September 1994), no intense precipitation was observed over the area of northern Italy, where the operational high-resolution day-5 forecast had predicted significant amount of precipitation.

4. Prediction of the Italian storm

The area of interest for all three cases is taken to be the greater Mediterranean region (GMR) with latitude between 30 and 50 degrees north, and longitude between 10 degrees west and 40 degrees east. In all three cases we concentrate on the ECMWF T213L31 and EPS predictions

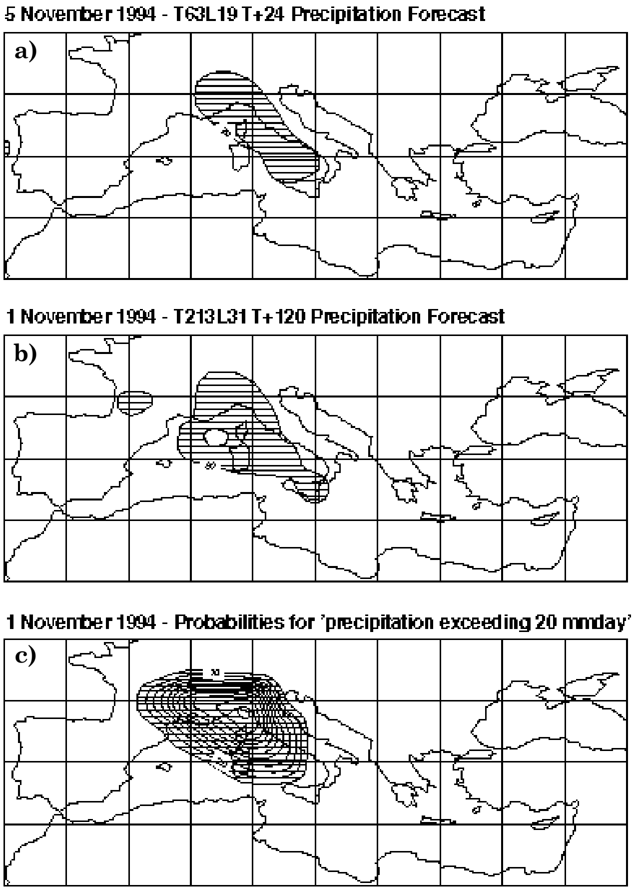

***Fig. 2:** a) T63L19 precipitation prediction of the forecast initiated on 94.11.05, accumulated between day 0 and day 1 (areas over 20 mm/day are shaded). b) T213L31 precipitation prediction of the forecast initiated on 94.11.01, accumulated between day 4 and day 5 (areas over 20 mm/day are shaded). c) EPS probability values for the event 'total precipitation exceeding 20 mm/day' for the EPS started on 94.11.01 (precipitation accumulated between day 4 and day 5). Contour isolines every 10%.*

5 days prior to the event. The forecast of the accumulated precipitation between 12 UTC of the 5th and the 6th of November 1994 (day-5 precipitation forecast initiated on 01.11.94) is verified.

A detailed picture of the total precipitation observed over the area of interest between 12 UTC of the 5th and the 6th of November 1994 can be found in Buzzi et al, (1995). For this interval, values consistently exceeding 100 mm/day were reported. Here (and for the other two remaining cases) we consider the 24h precipitation forecast from T63 control to highlight areas where the observed precipitation exceeded the 20 mm/day (Fig. 2a). These areas agree well with areas obtained from a similar set of analyses made using direct (observed) rainfall.

The T213L31 day-5 precipitation forecast is very skilful (Fig. 2b). The predicted pattern is positioned almost correctly (as indicated by the shaded areas corresponding to values exceeding the 20 mm/day), and the rainfall amounts (a peak value of 135 mm/day is forecasted) give a clear warning of severe weather over the region of interest.

Figure 2c shows the day 4 to day 5 probability forecast initiated on 01.11.94 for the event 'total precipitation exceeding 20 mm/day' to occur. Note that a 10% probability means that about 3 ensemble members support the occurrence of the event. Also note that since the EPS resolution is T63L19, the relatively poor resolution can be a substantial handicap especially in orographic regions. However, overall the probability map strongly supports the T213L31 prediction, and thus provides the forecaster with additional confidence in the high-resolution forecast.

Additional confidence can be obtained by studying the consistency of the EPS from previous days. Figure 3 shows the evolution of the 2-dimensional probability maps from the EPS started 8, 7, .., 3 days before the event, all verifying over the same time interval. Significant possibility of an extreme event is already detectable in the day-6 forecast from the EPS initiated on 31.10.94 (Fig. 3c). Starting from the EPS initiated on 01.11.94, the signal of an extreme event becomes clear and consistent (Fig. 3d-f). The strong consistency between the three more recent EPS forecasts should add confidence in the T213L31 forecast.

Forecasters are frequently asked to issue local predictions together with forecasts valid for larger areas. Figure 4 shows the 10-day precipitation prediction for the T63 grid-point (landpoint) closest to the Piedmont area (with coordinates: lat 43.8 N - lon 7.5 E), given by the EPS started 8, 7, .., 3 days before the event (panels a, b, .., f, respectively of Figure 4). Note that values are reported every 12 hours, and thus represent total precipitation accumulated over a 12 hour time interval. A mark (dot) on each panel identifies the 12 hour accumulated precipitation from observations (obtained by averaging four observations surrounding the selected point, as being reported by Buzzi et al, 1995) valid for 00 UTC 06.11.94. Note how the dispersion of the predicted rainfall decreases as the EPS starting date gets closer to the verifying time interval, and how this corresponds to a better agreement between the predicted and observed value.

5. Prediction of the Hellenic storm

As in the previous case (Section 4) we concentrate on the operational T213L31 and EPS predictions 5 days prior to the event. The precipitation is verified from the observed accumulated precipitation between 12 UTC of the 21st and the 22nd of October 1994.

A detailed picture of the total precipitation observed over the area of interest between 12 UTC of the 21st and 22nd can be found in Lagouvardos et al, (1995). For the same reasons as discussed for the Italian case we use the 24h precipitation forecast from the T63 control to highlight areas where precipitation exceeds 20 mm/day (Fig. 5a).

Figure 5b shows the day-5 T213L31 forecast originating from 17.10.94. The operational day-5 forecast positions the precipitation maxima correctly (over the centre of the Hellenic pensisula), reaching a maximum of 125 mm/day (a clear indication that an extreme event is

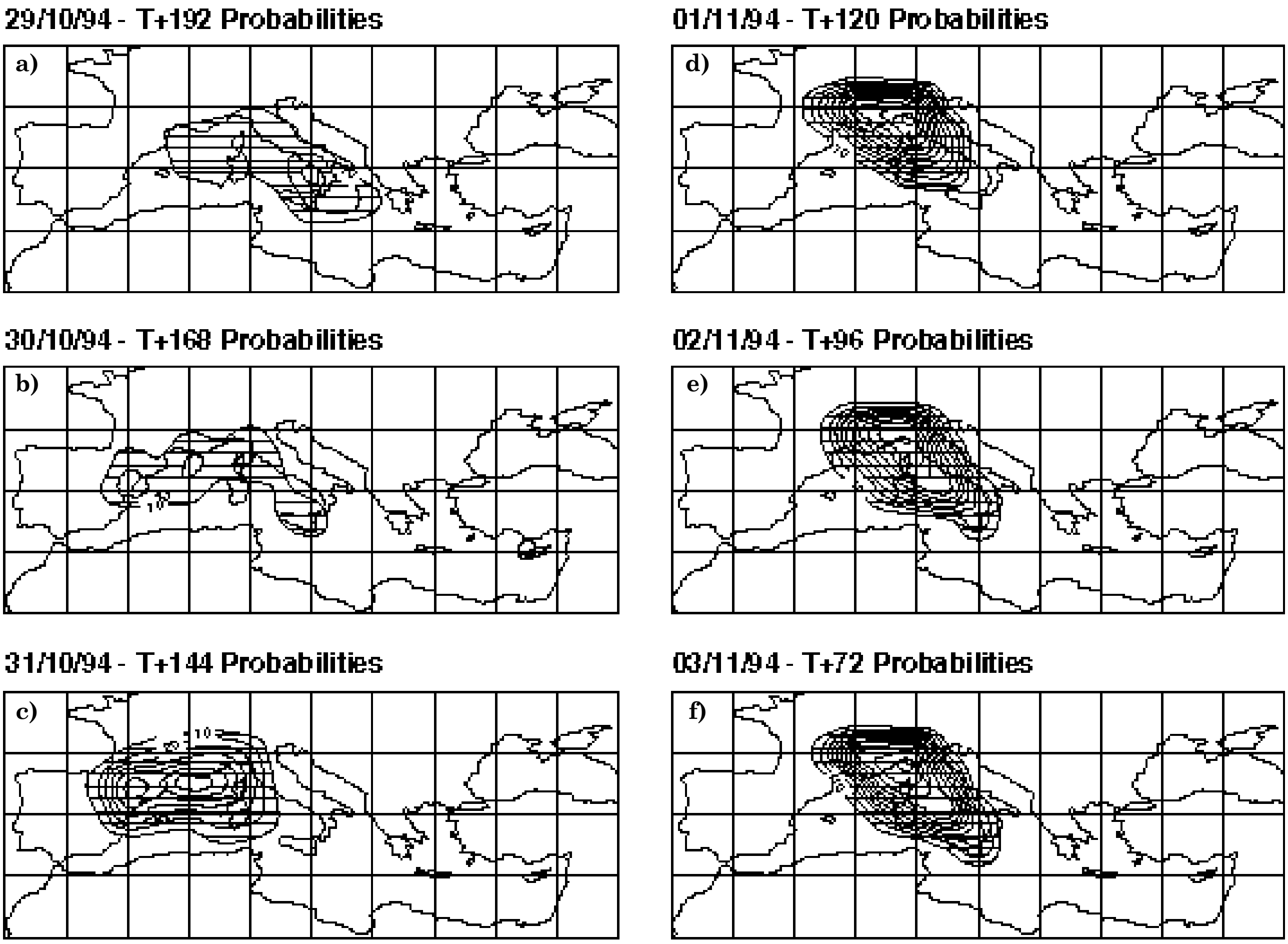

Fig. 3: *Probability prediction for the event 'total precipitation exceeding 20 mm/day' to occur between the 12 UTC of the 5th and the 6th of November 1994, for the EPS initiated on a) 94.10.29 (precipitation accumulated between day 7 and day 8), b) 94.10.30 (precipitation accumulated between day 6 and day 7), c) 94.10.31 (precipitation accumulated between day 5 and day 6), d) 94.11.01 (as Fig. 3c), e) 94.11.02 (precipitation accumulated between day 3 and day 4), and f) 94.11.02 (precipitation accumulated between day 2 and day 3). Contour isolines every 10%.*

likely to take place). There are suggestions of two other precipitation maxima the first over the southern Adriatic sea and the second over north Tunisia, but their intensity is much lower.

In Figure 5c the day-5 probability values for the event 'total precipitation exceeding 20 mm/day' to take place are shown, as predicted by the EPS initiated on 17.10.94 (i.e., constructed using the total precipitation ensemble member predictions between day 4 and day 5). The day-5 probability map reveals a high chance of an extreme event. Although the area of maximum probability (values reaching a value of 70%) is shifted over the northern Ionion sea (western sea area of Greece), the area of the main impact shows high probability values. It is clear that the EPS as it is expressed by the 2-D probability maps supports the operational precipitation forecast.

The evolution of the probability maps from the EPS initiated 8, 7, ..., 3 days prior the event can be seen in Figure 6. It is evident that a significant warning of an extreme event is becoming visible from day 6 (although shifted over the southern Italy - west Ionion sea) for the ensemble originated from 16.10.94 (Fig. 6c). In the previous two days low probability values are present over the wrong area (Figs. 6a-b). For the next three consecutive probability forecasts (forecast day-5, -4 and -3) the probabilities become consistent with the area of the main impact gradually shifting to the east and to a more correct position (Figs. 6d-f). It should be noted that the area of absolute maxima of probability is placed to the west of the main impact area. However, overall the EPS members support reasonably well the operational forecast.

6. The false-alarm case

We try to assess the information provided from EPS for the false-alarm case comparing the day-5 T213L31 deterministic forecast with the corresponding precipitation probability maps. We also study the evolution of the probability maps for the period 8 to 3 days prior the false alarm event.

In Figure 7a we use the 24h precipitation forecast from the T63 control run to highlight areas where precipitation exceeded the 20 mm/day (in well agreement with the corresponding observations of total precipitation for the same area and period).

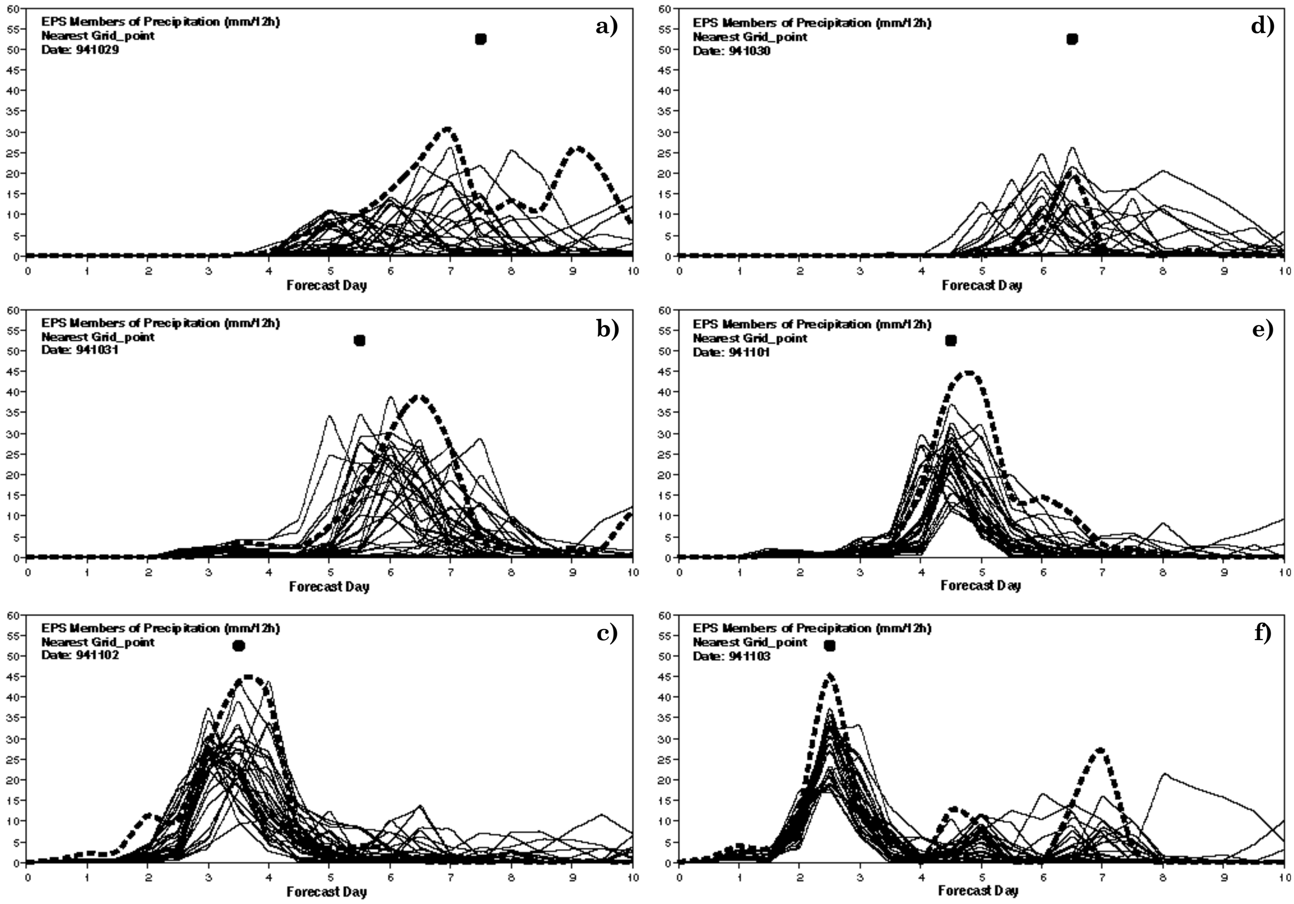

***Fig. 4:** 10-day EPS plumes of the T63 land grid-point closest to the Piedmont area (lat: 43.8 N - lon: 7.5 E), of total precipitation predicted by the T213L31 (bold dash) and by the 33 EPS members (thin solid), for the predictions started on a) 94.10.29, b) 94.10.30, c) 94.10.31, d) 94.11.01, e) 94.11.02 and f) 94.11.03. Values are reported every 12 hours. For each time t the corresponding value represents the precipitation cumulated between t-12h and t. A mark on each panel identifies the amount of the 12 hour accumulated precipitation from observations valid for 00 UTC 94.11.06.*

The 120 hour operational T213L31 precipitation prediction (total precipitation between day 4 and day 5) verifying for the time interval between 12 UTC of the 15th and the 16th of September, is shown in Figure 7b. The operational day-5 precipitation forecast originating from 12 UTC of 11.09.94 is considered to be poor. The operational precipitation forecast maxima (reaching a value of 65 mm/day) is wrongly positioned over northeastern Italy.

In Figure 7c the 120 hour probability values (two-dimensional probability maps) are shown for the event 'precipitation exceeding 20 mm/day' using precipitation fields from the EPS started on 11 September 1994 12 UTC. The day-5 probabilities do not support the operational T231L31 forecast at all. Contour interval of only 5% (half of the contour interval used in all other probability maps presented in this paper) had to be used to see at least one isoline (i.e. the isopleth of 5%).

The same interval (5%) is used in the next figure (Fig. 8) where the evolution of probability maps from the EPS initiated 8, 7, ..., 3 days prior the event is shown. All forecasts verify between the interval 12 UTC of the 15th and 16th of September 1994. It is evident that the EPS probability values are minimal over the area of interest (i.e. main impact area as suggested by the operational day-5 precipitation forecast) and they do not provide any support to the operational forecast throughout the whole interval.

7. Conclusions

In this paper, the ability of the ECMWF EPS to provide a measure of confidence in the high resolution operational forecast, has been studied. Three specific case studies were presented, corresponding to synoptic situations in which the high-resolution operational deterministic forecast predicted heavy rainfall in the medium range. In the first two cases, the EPS supported the high resolution model, and the deterministic forecast verified correctly. In the third case, the EPS did not support the high resolution model forecast, and this forecast did not verify correctly.

EPS forecasts of extreme weather events are necessarily compromised by the moderate resolution of the T63L19 model used to generate the ensembles. In future studies, ensembles will be made using at least T106L31 resolution. The benefits of such increases in model resolution will be compared with the benefits of increases in ensemble size. However, the ability of the higher resolution model to simulate heavy rainfall, strong winds, and extreme temperatures, suggests that high resolution will be necessary for the production of reliable probability forecasts of extreme weather.

Fig. 5 (right): *a) T63L19 precipitation prediction of the forecast initiated on 94.10.21, accumulated between day 0 and day 1 (areas over 20 mm/day are shaded). b) T213L31 precipitation prediction of the forecast initiated on 94.10.16, accumulated between day 4 and day 5 (areas over 20 mm/day are shaded). c) EPS probability values for the event 'total precipitation exceeding 20 mm/day' for the EPS started on 94.10.16 (precipitation accumulated between day 4 and day 5). Contour isolines every 10%.*

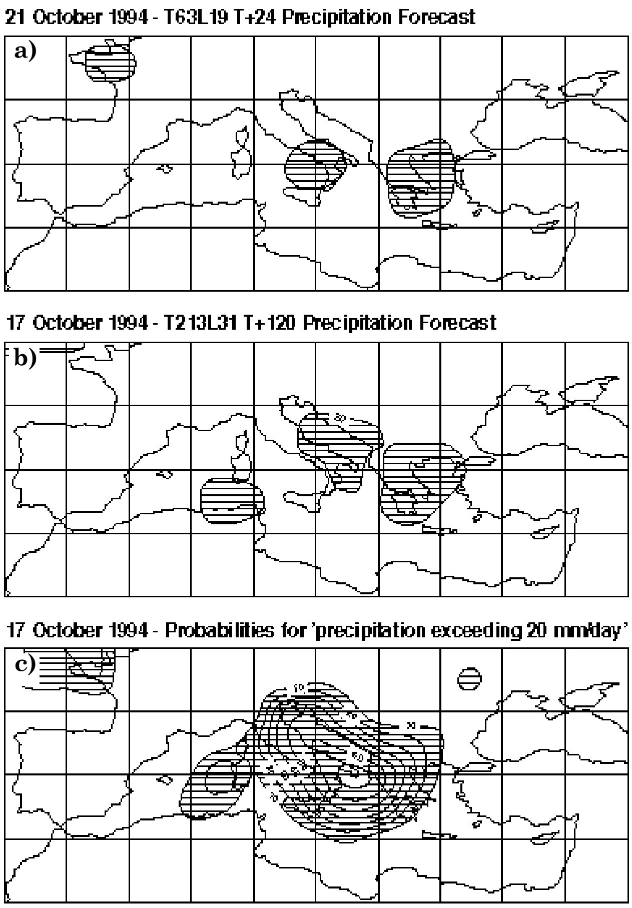

Fig. 6 (below): *Probability prediction for the event 'total precipitation exceeding 20 mm/day' to occur between the 12 UTC of the 21st and the 22nd of October 1994, for the EPS initiated on a) 94.10.14 (precipitation accumulated between day 7 and day 8), b) 94.10.15 (precipitation accumulated between day 6 and day 7), c) 94.10.16 (precipitation accumulated between day 5 and day 6), d) 94.10.17 (as Fig. 8c), e) 94.10.18 (precipitation accumulated between day 3 and day 4), and f) 94.10.19 (precipitation accumulated between day 2 and day 3). Contour isolines every 10%.*

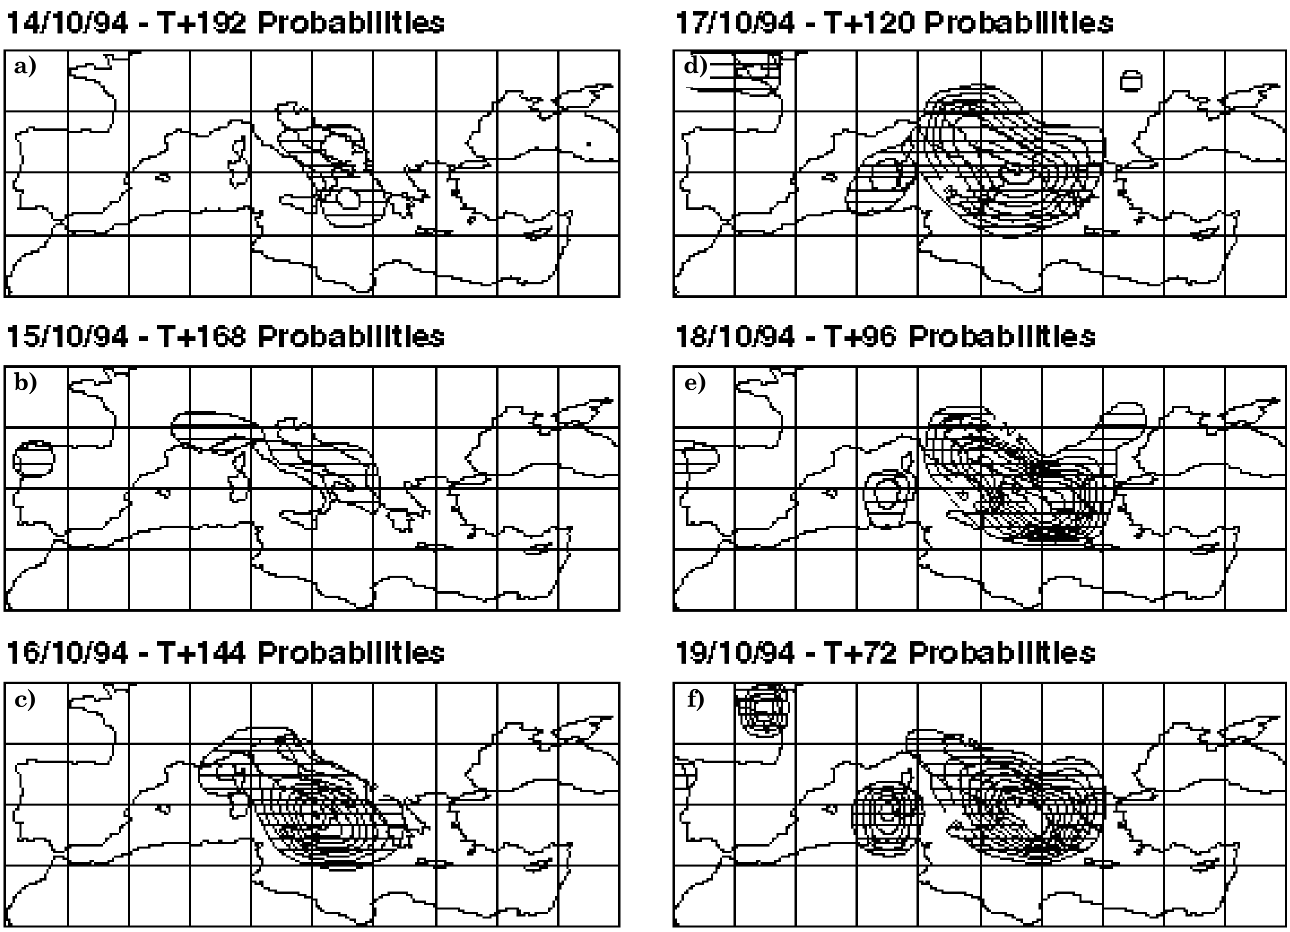

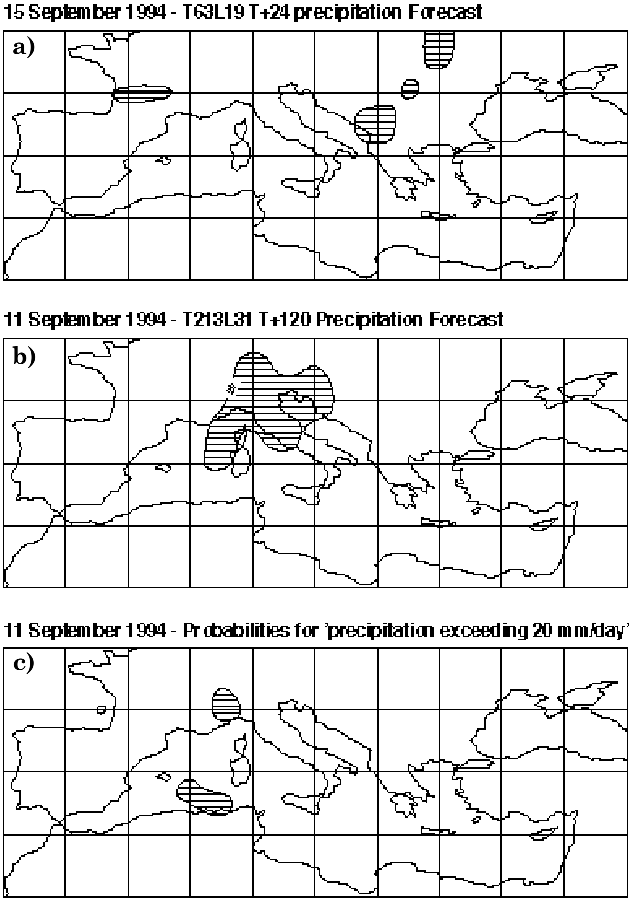

Fig. 7 (left): *a) T63L19 precipitation prediction of the forecast initiated on 94.09.15, accumulated between day 0 and day 1 (areas over 20 mm/day are shaded). b) T213L31 precipitation prediction of the forecast initiated on 94.09.11, accumulated between day 4 and day 5 (areas over 20 mm/day are shaded). c) EPS probability values for the event 'total precipitation exceeding 20 mm/day' for the EPS started on 94.09.11 (precipitation accumulated between day 4 and day 5). Contour isolines every 5%.*

Fig. 8 (below): *Probability prediction for the event 'total precipitation exceeding 20 mm/day' to occur between the 12 UTC of the 15th and the 16th of September 1994, for the EPS initiated on a) 94.09.08 (precipitation accumulated between day 7 and day 8), b) 94.09.09 (precipitation accumulated between day 6 and day 7), c) 94.09.10 (precipitation accumulated between day 5 and day 6), d) 94.09.11 (as Fig. 10c), e) 94.09.12 (precipitation accumulated between day 3 and day 4), and f) 94.09.13 (precipitation accumulated between day 2 and day 3). Contour isolines every 5%.*

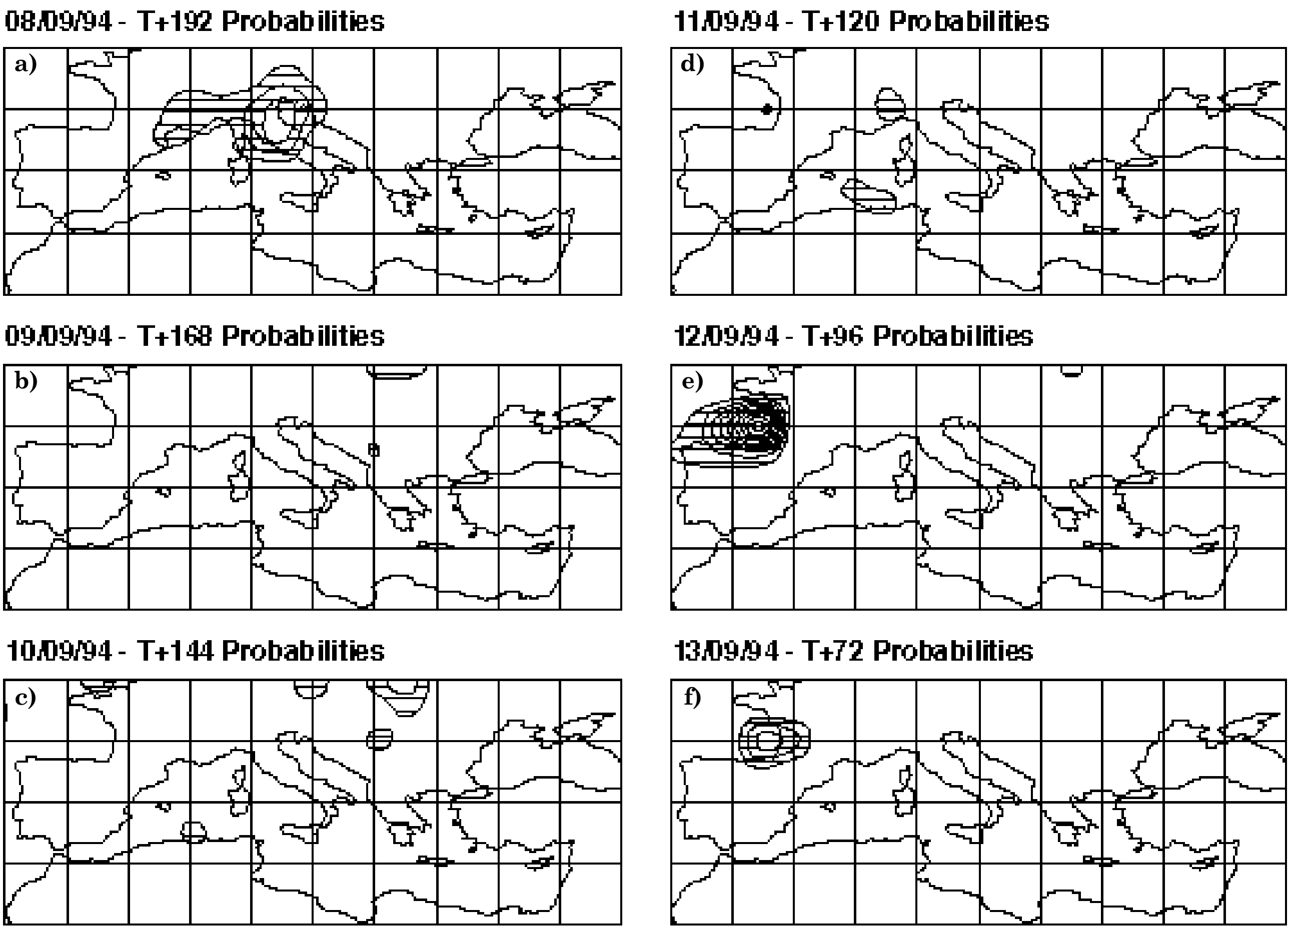

References

Athanassiadou M and A J Thorpe, 1995. The ANOMALIA project and MAP. *MAP Newsletter*, **3**, 48-50.

Binder, P and A Rossa, 1995. The Piedmont flood: Operational prediction by the Swiss Model. *MAP Newsletter*, **2**, 12-16.

Buizza, R and T N Palmer, 1995. The singular vector structure of the atmospheric general circulation. *J. Atmos. Sci.*, 52, **9**, 1647-1681.

Buzzi, A, C Cacciamani, T Paccagnella, P Patruno and N Tartaglione, 1995. Preliminary meteorological analysis of the Piedmont flood of November 1994. *MAP Newsletter*, **2**, 2-6.

Buzzi A and N Tartaglione, 1995. Meteorological modelling aspects of the Piedmont 1994 flood. *MAP Newsletter*, **3**, 27-28.

Courtier, P, C Freydier, J F Geleyn, F Rabier and M Rochas, 1991. The Arpege project at Meteo France. Proceedings of ECMWF Seminar on Numerical methods in atmospheric models, Shinfield Park, Reading RG2-9AX, UK, 9-13 September 1991, vol. 2, 193-231.

Dorninger M and M Hantel, 1995. Diagnosis of convection for the Piedmont flood of November 1994. *MAP Newsletter*, **3**, 40-42.

Epstein, E S, 1969. Stochastic dynamic prediction. *Tellus*, **21**, 739-759.

Fleming, R J, 1971a. On stochastic dynamic prediction. I: the energetics of uncertainty and the question of closure. *Mon. Wea. Rev.*, **99**, 851-872.

Fleming, R J, 1971b. On stochastic dynamic prediction. II: predictability and utility. *Mon. Wea. Rev.*, **99**, 927.938.

Gleeson, T A, 1970. Statistical-dynamical predictions. *J. Appl. Meteorol.*, **9**, 333-344.

Lagouvardos, K., V. Kotroni, S. Dobrivic, S. Nickovic and G. Kallos, 1995. The Storm of 2122 October 1994 over Greece: Observations and Model Results. Submitted to *J. G. R.*

Leith, C E, 1974. Theoretical skill of Montecarlo forecasts. *Mon. Wea. Rev.*, **102**, 409-418.

Molteni F, R Buizza, T N Palmer and T Petroliagis, 1996. The ECMWF Ensemble Prediction System: methodology and validation. *Q. J. R. Meteorol. Soc.*, **122**, 73-119.

Simmons, A J, D M Burridge, M Jarraud, C Girard and W Wergen, 1989. The ECMWF medium-range prediction models: development of the numerical formulations and the impact of increased resolution. *Meteorol. Atmos. Phys.*, **40**, 28-60.

T Petroliagis, R Buizza, A Lanzinger and T N Palmer

Summary of ECMWF Technical Report No. 77

ECMWF forecasts of the floods of January 1995

A. Lanzinger

The performance of the ECMWF operational T213 forecast model and the Ensemble Prediction System (EPS) during the period of widespread flooding over western Europe in January 1995 is reviewed in this ECMWF Technical Report.

During the second half of January 1995, large areas of western and central Europe received more than 100 mm of precipitation, with peak values of more than 200 mm. This resulted in widespread flooding, particularly over Brittany and the lower Rhine and Meuse regions.

The operational T213 forecast model showed consistently excellent performance over this period, in terms of forecasting the synoptic evolution and specifically the associated precipitation. At the 3-day range, the onset of the wet period, the precipitation amounts over these areas, and the timing of wet and intermediate dry periods were accurately predicted. In the range of five days there were difficulties with the onset around 17 January over western France as well as inaccuracies in the timing during the following period. However, precipitation amounts accumulated over several days were accurately predicted. The 7-day range still gave good guidance as to the synoptic regime and the main periods and magnitudes of intensity of precipitation over the area.

The accumulation of precipitation over several days of forecast, or even over the entire 10-day range, indicated well the amounts that were observed, with good estimation of peak values, but in some cases a shift of highest accumulations to the southwest. Associated high values of model soil wetness and runoff, although not strictly verified, indicate good performance of the surface parametrization scheme, at least in qualitative terms, and the potential usefulness for hydrological prediction of, for example, runoff.

The Ensemble Prediction Scheme also performed well during this period. Large scale flow patterns were correctly indicted in the medium range, with large majorities of the members of the ensemble supporting the high-resolution T213 forecasts. Consistently high probabilities of strong precipitation complemented the T213 forecast guidance.

Summary of ECMWF Technical Memorandum 218

A new sub-grid scale orographic drag parametrization: its formulation and testing

M Miller

In many of the studies concerning the representation of orography in Numerical Weather Prediction (NWP) and General Circulation (GCM) models, attention has been focused either on the parametrization of sub-grid scale mountain waves or on the optimal representation of the resolved mountain ranges. The first approach has led to the introduction of the gravity wave drag schemes (*Boer et al*, 1984; *Palmer et al*, 1986; *Miller et al*, 1989) and the second to the use of an envelope orography, for example, which improves the model representation of the large scale planetary waves (*Wallace et al*, 1983). These studies dealt essentially with the impact of the sub-grid scale orography and of the resolved scale orography on the global dynamics of the atmosphere. Recent studies of the local behaviour of the ECMWF model near the Pyrenees (*Lott*, 1995) have shown that the model underestimates the mountain drag, and generates mountain waves, with a horizontal scale close to the model truncation and which are often not observed. Furthermore, the way these waves are dissipated and affect the flow is unclear and unrealistic. These results are surprising since envelope orography increases the height and volume of the orography, and increases the drag on the atmosphere (*Tibaldi*, 1986). The envelope orography, however, is detrimental to the data assimilation process since more low level data, located below the model ground, are rejected using an envelope orography. Envelope orography also tends to give excessive precipitation, especially from convection generated by the elevated heating of the enhanced orography. It may therefore be desirable to replace the envelope orography by a mean orography without further changes but such a reduction of the mountain height will reduce the model mountain pressure drag with detrimental effect on the forecasts and model climate. The fact that the present gravity wave drag scheme inadequately represents low level drag, and the fact that the total mountain drag is already too small in the current model with envelope orography indicate that a major revision of the representation of Sub-grid Scale Orography (SSO) is desirable, and necessary to successfully represent the overall impact of orography on the global dynamics. *Clark and Miller* (1991) have shown, using a nested high resolution model, that there is a large underestimation of the total drag at horizontal resolutions coarser than about 10 km, which cannot adequately be made up by the use of an envelope orography.

Recent work by *Baines and Palmer* (1990) presents the principles of a sub-grid scale orographic drag scheme in which particular emphasis is placed on the representation of three-dimensional wave surface stress. These authors also suggested that further drag should be provided at model levels which intersect the sub-grid scale orography. ECMWF Technical Memorandum 218 (*F Lott and M Miller*, 1995) proposes a theoretical formulation for such a drag on model levels and this forms a major component of the new sub-grid scale orographic drag scheme. The general principles are presented, interpreting results from theoretical studies of flow near mesoscale orography in the context of numerical weather prediction models. The parametrization scheme is described in detail and, using an off-line procedure, the new scheme is used to predict the drag and the momentum flux profiles, observed during the PYREX experiment. The scheme is tested with the ECMWF model (at T106 and T213 resolutions), using forecast experiments covering the Periods of Intense Observation of the PYREX campaign. At both resolutions, it appears that the model with the new scheme is consistently able to reproduce the measured pressure drag. Isentropic diagnostics of the flow dynamics are used to study the impact of the scheme on the low level flow dynamics. Finally a summary is given of some experimental forecast results comparing performances of the model with mean orography with and without the new scheme, and with a version of the model with envelope orography and the old gravity wave drag scheme that has been used operationally at ECMWF.

In this new parametrization of sub-grid scale orographic drag, particular attention has been paid to the drag in the model layers of the atmosphere below the sub-grid scale mountain peaks. This low level part of the scheme replaces the envelope enhancement of model orography. The upper part of the scheme still represents the role of gravity waves. This part has been revised to allow a better representation of mountain ridge orientation and anisotropy. It also removes some arbitrary assumptions in the previous scheme. The depth over which the gravity wave drag part of the scheme is felt by the flow at low levels is now linked to dynamical properties of the mountain waves. The incident wind characteristics are calculated above the boundary layer (when the mountain is high enough) suppressing the strong diurnal cycle which was found to occur in the old gravity wave drag scheme. The parameters of the new scheme have been calibrated using an off line procedure, in which the scheme has been used to predict the mountain drag measured during PYREX. In this calibration, the new drag scheme outperforms the old gravity wave drag scheme, and fits well with the surface drag measured during field experiments. It also gives more realistic momentum flux vertical profiles than those of the old scheme. It was further shown that the new scheme has a realistic dynamical impact on the model dynamics in the vicinity of mountains, reinforcing the scheme's basic premise that mesoscale mountain drag slows down the low-level flow under most atmospheric situations. Results from forecast experimentation such as skill scores, precipitation amounts and distribution indicate that overall the combination of mean orography together with the new scheme performed better than that using envelope orography plus the current GWD scheme. However, there remains the problem of representing the lateral force present when there is strong flow parallel to the ridges. This is the subject of ongoing work. The scheme is beneficial to the forecast skill of the ECMWF model at all forecast ranges and has been in operational use since 4 April 1995[1]. The theoretical background to this work suggests that not only models of climate and GCM resolution but also much higher resolution limited area models should parametrize the drag due to the 'blocked' flow.

[1] See ECMWF Newsletter 70, pp 2-8.

References:

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

Boer, G J, N A Mc Farlane, R Laprise, J D Henderson and J P Blanchet, 1984: The Canadian Climate Centre spectral atmospheric general circulation model. Atmos-Ocean, **22**, pp 397-429.

Clark, T L and M J Miller, 1991: Pressure drag and momentum fluxes due to the Alps. II Representation in large-scale atmospheric models. *Q J Roy Meteor Soc*, **117**, 527-552.

Lott, F, 1995: Comparison between the orographic response of the ECMWF model and the PYREX 1990 data. To appear in *Quart J Roy Met Soc.*

Lott, F and M J Miller, 1995: A new sub-grid scale orographic drag parametrization: its formulation and testing. *ECMWF Technical Memorandum 218* and *Quart. J Roy. Met. Soc.,* **123** (1997)

Miller, M J, T N Palmer and R Swinbank, 1989: Parametrization and influence of sub-grid scale orography in general circulation and numerical weather prediction models. *Meteor & Atmos Phys*, **40**, 84-109.

Palmer, T N, G J Shutts and R Swinbank, 1986: Alleviation of systematic westerly bias in general circulation and numerical weather prediction models through an orographic gravity wave drag parametrization. *Quart J Roy Met Soc*, **112**, pp 2056-2066.

Tibaldi, S, 1986: Envelope orography and maintenance of quasi-stationary waves in the ECMWF model. *Adv Geophys*, **29**, 339-374.

Wallace, J M, S Tibaldi and A Simmons, 1983: Reduction of systematic forecast errors in the ECMWF model through the introduction of an envelope orography. *Quart J Roy Met Soc*, **109**, pp 683-717.

Distributed memory computer systems at ECMWF

ECMWF has had a long history of developing and using parallel processing methods. Initially, this was done on shared memory systems, more recently on distributed memory systems. Previous Newsletter articles have covered some of the shared memory system work (e.g. Newsletter No. 60, December 1992), this article reviews the history, current situation, and future plans of ECMWF's involvement with distributed memory systems. It is based on a talk given at the Supercomputer '96 meeting in Mannheim (20-22 June 1996).

1. Introduction

The computing requirements of numerical weather prediction (NWP) necessitate the exploitation of all available computing resources simultaneously for one task. Therefore, ECMWF started to use parallel processing in 1983, as soon as it became available commercially for high performance computing. For the same reason, ECMWF began preparations to use highly parallel systems with a distributed memory architecture well before acquiring such a system in 1996. The different steps taken to achieve this position are described together with its challenges. Some results from benchmark runs are presented. Finally, an outlook towards the operational use of distributed memory systems with a large number of processors is given.

Due to the operational requirement that its medium range global weather forecasts are completed within a tight schedule, and due to the computing power required for atmospheric modelling (see [3]), ECMWF has had to concentrate from the beginning on the use of high-speed computers. Initially, this led to the study and exploitation of vector processing when ECMWF received the first European CRAY-1A in 1978. Later this was complemented by the use of parallelism or *"multi-tasking"*, after the CRAY-1A was replaced by a CRAY X-MP/22 in 1983.

In addition, ECMWF recognized relatively early on that the future requirements of meteorological modelling could only be met by massively parallel systems. Therefore in 1984, it started a series of biennial workshops on *'The Use Of Parallel Processors In Meteorology'.*

There were some early experiences in the 1970's with parallel systems for numerical weather prediction at FNMOC (USA), when a Control Data 6500 with two processors was used to calculate both hemispheres in parallel and to exchange the boundaries via ECS memory. However, it took the advent of the CRAY X-MP series in 1982 before the parallelism inherent in numerical weather prediction was exploited again in a wider context. The publications by Dent between 1984 and 1988 (see [5], [6], [7]) provide an insight into the techniques used and results gained by the evolving parallelising strategies. Using "macrotasking", it was possible to foresee the efficient use of shared memory systems with of the order of 100 processors.

ECMWF used the spectral model described by Dent in its various forms from 1983 to 1994 on a sequence of shared memory vector processors built by Cray Research, Inc. They spanned the range from an X-MP/22 to a C90/16-2048, with an X-MP/48 and a Y-MP864 in between. The performance of the model on the different platforms has been good, finally achieving an average of 5.5 GFlops/s for the duration of a complete 10-day forecast run on the CRAY C90/16-2048, i.e. more than 35% of its peak.

As outlined in [13], the current developments in operational analysis, high-resolution forecasting and ensemble prediction systems require a major increase in computing resources compared to that which is available today from a CRAY C-90/16. For example, by the year 1999, the computing speed will have to be raised by a factor of more than 20 over this machine to allow the operational implementation of an effective four-dimensional variational data assimilation system.

Assuming that no revolutionary breakthrough in technology occurs, the number of processors in a shared memory configuration is probably restricted to less than 100 due to physical memory access limitations. Thus an upper limit for the possible performance of such systems is set. Therefore, only massively parallel computer systems with a physically distributed memory architecture can at present be envisaged to reach the performance levels required by NWP.

2. Preparation for distributed memory systems

2.1 Biennial workshops on "Use of Parallel Processors in Meteorology"

Even without the explicit requirement for performance levels that necessitate distributed memory systems as outlined above, it had become clear already in 1984 that only massively parallel systems would in the longer term be able to meet the computing requirements of NWP. It was therefore decided to canvass the views of the researchers in the fields of NWP, Computer Science and Numerical Mathematics and to expose their findings to the scrutiny of the computer vendors. Thus, in 1984, the idea of holding biennial workshops at ECMWF was born and the first of the series was held at ECMWF in autumn 1984.

Initially, (1984 and 1986), the workshops concentrated on the question of parallelism in general and less on specific programming paradigms. For example, there was a lively discussion whether particular model representations, such as spectral or grid-point based, were better suited for parallel systems, and whether there existed some affinity between certain machine architectures and model formulations. In addition, the market for parallel systems was divided into protagonists for MIMD (Multiple Instruction Multiple Data) and SIMD (Single Instruction Multiple Data). There was not even a consensus as to whether MPP systems with O(100) processors would ever be efficient for running weather prediction models, due to problems foreseen in the degree of parallelization, synchronisation, high latency of message passing, I/O etc.

By 1988, the outlook for parallel processing became more optimistic. There were early examples of models running efficiently in parallel on O(10) processors, even though the models were still very simple, e.g. shallow water codes or models without physics. This development continued until 1992 when for the first time a variety of realistic models could be demonstrated running on parallel systems achieving acceptable levels of efficiency.

By 1994, SIMD systems had lost their appeal for the NWP community. It was a generally held belief that MIMD parallel machines with a distributed memory architecture would succeed the then current parallel shared memory vector systems in the not so distant future. There remained a number of challenges for the operational use of such systems, but the workshops had presented sufficient evidence to support the move to the new architecture.

With the mix of different expertise and viewpoints of the participants, the workshops have generated a lot of useful input for NWP and were instrumental in focusing the efforts of the various researchers in the area of parallelising atmospheric models. In addition, the close relationship with computer vendors meant that practical problems were not overlooked and that realistic benchmarks of systems could be carried out.

2.2 GENESIS project

Following the advice given at the first workshop of the above series, ECMWF joined the Esprit project GENESIS [9] in 1989 as an Associate Partner. Its goal for joining was "to define a good strategy, algorithmically and programming- wise, for introducing future massively parallel supercomputers into operational global weather forecasting" [12]. Working together with the partners of this project, ECMWF evaluated various strategies for parallelising spectral forecast models and experimented with different FORTRAN extensions to express the required parallel constructs. It used Helmholtz solvers and shallow water codes as its tools for benchmarking the available parallel computer systems.

The main outcome of ECMWF's participation in the two-year project was a good understanding of the parallelization strategies to be used for spectral models, and a decision on the programming paradigm that would be required to achieve portability between different parallel machines yet with little impact on performance.

The first result, the choice of the "transposition" strategy [1], proved very user friendly, because it concentrated all actions dealing with distributed memory access into a few isolated FORTRAN subroutines. All other routines then access their relevant data without any additional overhead, because all data is local to the processor working on them.

The second result, the selection of PARMACS [4] as the message passing interface for FORTRAN, enabled ECMWF to develop portable programs for benchmarking parallel systems. The possibility to provide a common code to all computer manufacturers for evaluating their offerings of parallel systems led to the foundation of RAPS.

2.3 RAPS

RAPS (Real Applications on Parallel Systems) was established by several European developers and users of large production codes[1] as an open consortium in 1992. It aims to create a realistic benchmark suite for parallel computer systems by achieving portability through the adoption of a joint programming model. RAPS cooperates with computer vendors in a Working Group[2]. It organizes regular open workshops and vendor specific seminars. The benchmark results of the RAPS codes are publically available.

The adoption of a RAPS programming model, initially FORTRAN 77 and PARMACS [11], then FORTRAN 90 and MPI from 1996, led to the acceptance of this programming paradigm by the associated computer vendors and made it possible for code developers to program in a reasonably portable manner (as long as I/O was left out).

The close co-operation between the developers of large parallel production codes and the representatives of the computer vendors has proven to be very successful for both sides. The developers were able to test their programs on a variety of hardware platforms and receive feed-back on performance issues, while the computer vendors could see and influence the performance characteristics of codes which would be used as benchmarks in later procurements.

A typical example of the impact that information received from the computer vendors had on the development of the code can be seen in [8]. The difference in code optimization for vector and cache supported RISC processors could be integrated in the benchmark code. In this way, a fairer comparison between vendors of different processor architecture could be achieved.

Benchmark results reported by Fujitsu as part of the RAPS activities during the workshop held at DWD on 15-16 November 1995 are shown in Table 1 below. They reflect the performance of the ECMWF IFS RAPS 3.0 submission [2] with real test data on Fujitsu's VPP500 systems. Scalability tests, mainly undertaken by Cray Research, confirmed that the IFS code could efficiently run on 1024 T3D processors achieving a speed-up of about 900.

[1] At present: CERFACS, CERN, DKRZ, DLR, DWD, ECMWF, GMD, METEOFRANCE, NERC, PALLAS GmbH, UKMO, University of Southampton

[2] At present: Convex, Cray Research, Fujitsu, Hitachi, IBM, NEC, SNI

	Nodes	Time	Speed-up	Efficiency	Factor to CRAY C90/16
46	1155	1	100.00%	3.9	
52	1015	1.14	100.01%	4.4	
104	571	2.02	99.88%	7.9	

Table 1: IFS T213L31 SL benchmark results on Fujitsu VPP 500

The participation of ECMWF in RAPS proved that IFS could run efficiently on O(1000) processors and that the IFS performance on distributed memory systems could significantly surpass that achievable on the most powerful shared memory system of the day.

2.4 Operational tests

In addition to the knowledge gained by having been a partner in GENESIS and the on-going information received from RAPS, ECMWF also needed operational experience with highly parallel systems before it could seriously consider moving its operational activities to distributed memory systems. It therefore joined the Esprit Project GP-MIMD2 (P7255) from 1 March 1993 to 31 October 1994 as an Associate Partner to evaluate a Meiko CS-2 system installed at CERFACS. Furthermore, ECMWF took delivery of a 128 processor CRAY T3D system in August 1994.

As a result of these activities, in particular the operational experience gained by having a distributed memory system installed and heavily used, ECMWF embarked towards the end of 1994 on the full scale migration of all its operational software to distributed memory systems using the portable RAPS programming model. The CRAY T3D served as a test-bed for the migration work.

2.5 Invitation to Tender (ITT)

Taking the foreseeable computing requirements mentioned above as a guideline, ECMWF then issued an open ITT in early 1995 to replace its CRAY C90 and T3D systems in 1996. The minimum required performance level was such that only distributed memory systems could attain it. ECMWF was therefore faced with a choice between competing parallel systems from Cray (T3E), Fujitsu (VPP), IBM (SP), Meiko (CS-2) and NEC (SX-4).

After a detailed evaluation, and with the agreement of its Member States, ECMWF signed a service contract on 7 December 1995 with Fujitsu to supply a series of VPP300 and VPP700 systems between 1 April 1996 and 31 December 2000. The number of processors will increase from 16 for the initial VPP300 in 1996 to over 200 for the final VPP700 in 1998.

3. Initial operational use of Fujitsu VPP systems

The first Fujitsu system to be delivered as part of the service contract, a VPP300/16, was installed on 15 March 1996 and passed its Preliminary Acceptance Test (PA) on 2 April 1996. The PA includes tests to prove that the system meets its functional specification, a performance test to verify the claims of the vendor, and an error-free operation of the system over 24 hours. The PA was followed by a reliability test during which the system had to achieve, in normal operations, contractually defined levels of availability. This test was successfully passed on 14 July.

Following the PA, the system is being used to migrate all operational programs and to acquaint the registered users of ECMWF with the new system. In addition, the operational staff have to set up procedures for the operators, find the most suitable set of system parameters for ECMWF's workload, and to gain insight into its particular features.

On 25 June 1996, the second Fujitsu system, a VPP700/46, arrived at ECMWF to start its set of acceptance tests. It passed its PA test on 2 July and was then handed to ECMWF staff to set up its operating environment. On 22 July, a trial service for all users was started.

The operational workload was moved from the CRAY C90 to the VPP700 by September 1996. At that time, all remaining migration of programs had been completed and an initial set-up of the system implemented. The Cray service was terminated on 30 September 1996.

4. Outlook

After the successful move of operations from a parallel shared memory environment to a distributed memory system, a number of longer term challenges will have to be addressed (see [10]). Just two are mentioned here.

♦ **Parallel I/O**

A standard way to express parallel I/O within FORTRAN 90 has to be defined, otherwise the portability of programs will suffer. Furthermore, it has to be possible to write/read in parallel at high speed from a number of processors, but also be able to access the data via nfs or ftp.

♦ **Parallelising compilers**

Even though the message passing paradigm for programming will remain the most efficient way of writing parallel programs for some time to come, the move of computer vendors towards distributed memory systems with shared memory nodes having multiple CPUs requires multi-level parallelism to be considered. Since it is highly impractical to implement such concepts using message passing constructs, the only viable solution is to use parallelising compilers for the node level and message passing for the intra-node distribution.

G.-R. Hoffmann

References

[1] **S.R.M. Barros, T. Kauranne**: Spectral and Multigrid Spherical Helmholtz Equation Solvers on Distributed Memory Parallel Computers. Proceedings of the Fourth ECMWF Workshop on the Use of Parallel Processors in Meteorology. ECMWF. 1992. pp. 1-25.

[2] **S.R.M. Barros, D. Dent, L. Isaksen, G. Robinson**: The IFS Model - Overview and Parallel Strategies. Coming Of Age. G.-R. Hoffmann, N. Kreitz (eds). 1995. ISBN 981-02-2211-4. pp. 303-318.

[3] **L. Bengtsson**: Computer Requirements for Atmospheric Modelling. Multiprocessing in Meteorological Models. G.-R. Hoffmann, D.F. Snelling (eds.). 1988. ISBN 3-540-18457-0. pp. 108-123.

[4] **L. Bomans, R. Hempel, D. Roose**: The Argonne/GMD macros in FORTRAN for portable parallel programming and their implementation on the Intel iPSC/2. Parallel Computing 15 (1990). pp. 119-132.

[5] **D. Dent**: The Multitasking Spectral Model at ECMWF. Multiprocessing in Meteorological Models. G.-R. Hoffmann, D.F. Snelling (eds.). 1988. ISBN 3-540-18457-0. pp. 203-213.

[6] **D. Dent**: The ECMWF Model: Past, Present and Future. Multiprocessing in Meteorological Models. G.-R. Hoffmann, D.F. Snelling (eds.). 1988. ISBN 3-540-18457-0. pp. 369-381.

[7] **D. Dent**: A Modestly Parallel Model. The Dawn of Massively Parallel Processing in Meteorology. G.-R. Hoffmann, D.K. Maretis (eds.). 1990. pp. 21-31.

[8] **D. Dent, L. Isaksen, G. Mozdzynski, G. Robinson, F. Wollenweber, M. O'Keefe**: The IFS Model Performance Measurements. Coming Of Age. G.-R. Hoffmann, N. Kreitz (eds). 1995. ISBN 981-02-2211-4. pp. 352-369.

[9] **T. Hey**: The Genesis Esprit project - An overview. Parallel Computing 20(1994). pp. 1605-1612.

[10] **G.-R. Hoffmann**: Areas of Concern for Multiprocessing - Revisited after 10 Years. Coming Of Age. G.-R. Hoffmann, N. Kreitz (eds). 1995. ISBN 981-02-2211-4. pp. 123-127.

[11] **H.-C. Hoppe**: The RAPS Programming Model 0. Parallel Supercomputing in Atmospheric Science. G.-R. Hoffmann, T. Kauranne (eds.). 1993. ISBN 981-02-1429-4. pp. 212-239.

[12] **T. Kauranne**: Summary of GENESIS work at the European Centre for Medium-range Weather Forecasts (ECMWF). Parallel Computing 20(1994). pp. 1685-1688.

[13] **A. J. Simmons**: High-Performance Computing Requirements for Medium-Range Weather Forecasting. Coming Of Age. G-R Hoffmann, N. Kreitz (eds). 1995. ISBN 981-02-2211-4. pp. 1-11.

Smart Card access to ECMWF's computers - an update

Two previous Newsletter articles have discussed why a smart card access system was considered necessary (Newsletter No. 67, pp. 27-33), and the initial implementation (Newsletter No. 70, pp. 1819). This article brings you up-to-date with current progress, the smart card controlled services now available, and future plans.

To summarise the first stages, in early 1994 the Centre conducted a "secure batch trial" to evaluate whether it would be possible to adapt a smart card system to authenticate users in a batch job environment. An experimental system was constructed and the principle of remote batch job authentication using smart cards was successfully demonstrated.

Based upon the results of this trial, the Centre proposed to the Technical Advisory Committee in September 1994 a strategy for secure access to the Centre's UNIX services. Consequently the Technical Advisory Committee recommended to Council that a smart card authentication system be introduced. This recommendation was accepted by the Council during its session in December 1994. SecurID smart cards from Security Dynamics were chosen as the method of control.

The implementation of Secure Batch has largely followed the proposed strategy, and in some areas additional security improvements have been achieved which were not foreseen in 1994. Telnet, batch and file transfer utilities are all in operation. They have been adopted by the Member States in varying degrees.

Telnet

Since 1 November 1995 the only UNIX server directly reachable from Member States via telnet has been 'ecgate1@ecmwf.int' ('ecserver@ecmwf.int' is an alias). All external access to it is authenticated via SecurID smart cards. From 'ecgate1@ecmwf.int' batch jobs can be submitted and a remote shell (rsh) invoked to the Centre's Fujitsu systems without further validation.

The difficulties originally foreseen in designing an administrative utility which would allow Member State Computing Representatives to assign spare cards or reset PIN codes for users in their country were overcome, and such a utility was put into operation in May 1995.

Batch (ecqsub)

Remote batch access using SecurID cards has been in operation since May 1996. The ecbatch utility allows Member State users to submit batch jobs to ECMWF (the 'ecqsub' function), and provides some job management functions (job status, job deletion, etc.). The ecbatch client software is available to all Member States and is supported on all major UNIX platforms.

To use the SecurID Batch System the user needs to be logged in on a Member State machine which has TCP/IP access to 'ecbatch@ecmwf.int', and which has the ECBATCH Client software installed. The user also requires a valid ECMWF SecurID Card in order to generate a certificate. To do this, the user will be prompted for his ECMWF user identifier and the PASSCODE. The certificate contains user information, a time stamp and the PASSCODE in PGP-encrypted form (PGP is a public domain encryption package).

When the Member State user enters ecqsub (or any other ECBATCH request) the ECBATCH client will check if there is a valid certificate. If there is no valid certificate, one will be generated. The ECBATCH client will then connect to the ECBATCH server running on an ECMWF host and transmit the job script, certificate and a checksum.

The certificate is valid for further ecqsub requests for a period of 12 hours from the time of its generation (7 days for ecqdel and ecqstat). The concept of a validity period for a certificate deviates from the original strategy, but has been deemed sufficiently secure given that the use of PGP encryption eliminates the possibility of an intruder monitoring network traffic, copying a certificate, and submitting a modified request. The PGP encryption is exclusively used to ensure the integrity of the data and authenticity of any request; it is not used for confidentiality, i.e. to encrypt the data.

File transfer (ftp, eccopy)

Incoming ftp requests require strong authentication. Thus 'ecgate1@ecmwf.int' now offers a version of ftp which prompts for smart card authentication rather than passwords. Data destined for another ECMWF host have to be transferred to 'ecgate1@ecmwf.int' first, and then copied to the other host or accessed via NFS.

Eccopy provides a means to send files from ECMWF to a Member State without the need for ftp. It is also used to return output from jobs submitted via the 'ecqsub' function of the ecbatch utility. For those Member States still without TCP/IP access, automated file transfers from ECMWF to Member States are now achieved via eccopy to the Centre's VAX cluster. This replaces the old 'sendtm' utility.

The 'eccopy' utility has been in operation since July 1995. It is available to all Member States and is supported on all major UNIX platforms.

MARS

MARS client software (which will shortly be available) also uses ECBATCH software certificates to validate MARS retrievals. This software, which will run on a Member State UNIX system, allows users to submit MARS retrieval requests direct to the MARS system at ECMWF, without the need for batch jobs or interactive sessions to be run at ECMWF. The user generates locally a MARS request, and if there is no valid certificate, the procedure to generate one will be started. The Member State MARS client will then connect to the MARS server running on an ECMWF host and transmit the request, certificate and a checksum. The certificate is validated with the ECBATCH server before the MARS request is serviced. Certificates for MARS requests expire after three months. This longevity makes MARS requests suitable for automated processing within Member States.

The validation of MARS requests from Cooperating Member States happens in the same fashion as for full Member States, but for them all other ECBATCH functions are disabled.

Future plans

Currently, eccopy allows only restricted file transfers from ECMWF to a Member State machine and is initiated from ECMWF. It is proposed to enhance the eccopy and the ecbatch utilities to allow file transfers in both directions between ECMWF and Member State machines. These file transfers could be initiated from ECMWF as well as from a Member State machine, and would require a valid certificate of similar lifetime as for the ecqsub function.

The ECBATCH software has currently no function to transfer data files, but an ecput and an ecget function may be added to allow file transfers between a Member State machine and ECMWF initiated from the Member State machine.

Finally it is planned to enhance the administrative utility for Member State Computer Representatives to allow them to modify the validity period of certificates for users in their country. This is currently a manual process involving an ECMWF system administrator.

Dieter Niebel

Fortran 95

The following article is reprinted from the CERN Computer Newsletter by kind permission of the author.

Why a new standard?

Following the publication of the Fortran 90 standard in 1991, two things happened. The first was that the High Performance Fortran Forum (HPFF) was set up to define a set of extensions to Fortran, to make it possible to write portable code when using parallel computers for handling problems involving large sets of data that can be represented by regular grids. This version of Fortran is known as High Performance Fortran (HPF), and is based on Fortran 90, mainly because of its array features. HPF is of a superset of Fortran 90, the main extensions being in the form of directives that take the form of Fortran 90 comment lines, and are thus recognized as directives only by an HPF processor. However, it did become necessary also to add some additional syntax, as not all the desired features could be accommodated in the form of such directives *(see also High Performance Fortran, CERN/CN/95/11)*.

The second happening was that the standards committees, X3J3 and WG5, decided on a strategy whereby a minor revision of Fortran 90 would be prepared by the mid-1990s - Fortran 95 - and a further revision by about the year 2000. At the same time, it was clearly desirable to include the new syntax defined by HPFF in Fortran 95 and, indeed, the HPF features are the most significant new features that it introduces.

The other changes consist mainly of what are known as 'corrections, clarifications and interpretations'. These came about as it was discovered that the text of the Fortran 90 standard contained a number of errors that required correction, some obscure wording that required further textual clarification, and ambiguous statements that required interpretation. Apart from the HPF syntax and the corrections, only a small number of other pressing but minor language changes were made. Fortran 95 is almost completely backwards compatible with Fortran 90.

A draft of the Fortran 95 standard was published in 1995, (ACM Fortran Forum, June, 1995). It was finalized in November 1995, and the expected schedule is for a new ISO standard, replacing Fortran 90, to be published in October 1996. This article is a summary of what is new in Fortran 95; the final version is completely described in Fortran 90/95 Explained, by John Reid and myself, (Oxford U. Press, March, 1996).

Features for parallel computing

This section describes new features for parallel computing, based on HPF. They are designed to make Fortran a still better language for applications on parallel architectures.

Pure procedures

Within FORALL statements or constructs (see below), the possibility that a function might have side effects is a severe impediment to optimization on a parallel processor - the order of execution of the assignments could affect the results. To control this situation, it is possible for the programmer to specify that a procedure has no side-effects by adding the PURE keyword to the SUBROUTINE or FUNCTION statement. In practical terms, this is an assertion that the procedure alters no global variable, performs no I/O, has no saved variables, and does not alter its INTENT(IN) arguments. These requirements are so formulated that a compiler can check that this is the case. An example is

```
PURE FUNCTION calculate (x)
```

All the intrinsic functions are pure.

Elemental procedures

Elemental procedures are those with scalar dummy arguments that may be called with array actual arguments provided that the array arguments have the same shape. For a function, the shape of the result is the shape of the array arguments. Many intrinsic functions are elemental and Fortran 95 extends this to non-intrinsic procedures. It requires an ELEMENTAL prefix on the FUNCTION or SUBROUTINE statement. This is very useful to the programmer who can get the same effect in Fortran 90 only by writing 22 versions, for ranks 0-0, 0-1, 1-1, 0-2, 2-0, 2-2, ... 7-7, and is an aid to optimization on parallel processors. An elemental procedure must satisfy all the requirements of a pure procedure, as it automatically has the PURE attribute. All dummy arguments and function results must be scalar variables without the pointer attribute. For an elemental subroutine, if any argument is array valued, all the arguments with intent INOUT or OUT must be arrays. For example, we can make a subroutine, swap , perform its task on arrays of any shape or size:

```
ELEMENTAL SUBROUTINE swap(a, b)
   REAL, INTENT(INOUT)  :: a, b
   REAL                          work
   work = a
   a = b
   b = work
END SUBROUTINE swap
```

The FORALL statement and construct

When a DO construct is executed, the processor is required to perform each successive iteration in order and one after the other. This represents a potential impediment to optimization on a parallel processor so, for this purpose, Fortran 95 has the FORALL statement, for example

```
FORALL(i = 1:n) a(i, i) = x(i)
```

which specifies that the individual assignments may be carried out in any order, and even simultaneously. The FORALL statement may be considered to be an array assignment expressed with the help of indices. In this particular example, this operation could not otherwise be represented as a simple array assignment. Another example of the FORALL statement is

```
FORALL(i=1:n, j=1:n, y(i,j)/=0.) &
   x(j,i)=1.0/y(i,j)
```

where the masking condition ensures that the assignment is not carried out for zero elements of Y.

The FORALL construct also exists. It allows several assignment statements to be executed in order. The FORALL equivalent of the array assignments

```
a(2:n-1,2:n-1) = a(2:n-1,1:n-2) &
               + a(2:n-1,3:n)   &
               + a(1:n-2,2:n-1) &
               + a(3:n,2:n-1)
b(2:n-1,2:n-1) = a(2:n-1,2:n-1)
```

is

```
FORALL(i = 2:n-1, j = 2:n-1)
   a(i,j) = a(i,j-1) + a(i,j+1) &
          + a(i-1,j) + a(i+1,j)
   b(i,j) = a(i,j)
END FORALL
```

This sets each internal element of A equal to the sum of its four nearest neighbours and copies the result to B. The FORALL version is more readable. Each assignment in a FORALL is like an array assignment; the effect is as if all the expressions were evaluated in any order, held in temporary storage, then all the assignments performed in any order. The first statement must fully complete before the second can begin.

A FORALL statement or construct may contain pointer assignments, may be nested, and may include a WHERE statement or construct. Procedures may be referenced within a FORALL, both in the logical scalar expression that forms the optional mask or, directly or indirectly, in the body of the construct. All such procedures must be pure.

For full details of this important new feature, please see the references above.

WHERE construct extensions

It is permitted in Fortran 95 to mask not only the WHERE statement of the WHERE construct, but also any ELSEWHERE statement that it contains. A WHERE construct may contain any number of masked ELSEWHERE statements but at most one ELSEWHERE statement without a mask, and that must be the final one. In addition, WHERE constructs may be nested within one another, just like FORALL constructs. A WHERE assignment statement is permitted to be a defined assignment, provided that it is elemental. Finally, a WHERE construct may be named in the same way as other constructs.

Other new features

Improved compliance with IEEE arithmetic

Many processors distinguish at the hardware level between a negative real zero value and a positive real zero value, and the IEEE standard for binary floating-point arithmetic (IEEE 754 or IEC 559-1989) makes use of this where possible. In order to be able to distinguish between the two cases in Fortran 95, the function SIGN has been generalized such that the sign of the second argument may be taken into account even if its value is zero. On a processor that has IEEE arithmetic, the value of SIGN(2.0, -0.0) is -2.0.

Automatic deallocation of allocatable arrays

In Fortran 90, an allocated array that does not have the save attribute has an undefined allocation status on return from a subprogram. In Fortran 95, such an array is automatically deallocated. This feature not only avoids inadvertent memory leakage, but prevents the very undesirable undefined allocation status. For allocatable arrays defined in modules, see the references.

Specification functions

Any of the intrinsic functions defined by the standard may now be used in a specification expression. In addition, a non-intrinsic pure function may be so used provided that certain simple conditions are fulfilled. Such functions are called specification functions. This feature will be a convenience for specification expressions that cannot be written as simple expressions.

Pointer initialization and the function NULL

Fortran 95 allows pointers to be given the initial status of disassociated in a type declaration statement such as

```
REAL, POINTER, DIMENSION(:) :: vector => NULL()
```

NULL is a new intrinsic function and may be used in an executable statement:

```
vector => NULL()
```

It is very useful in avoiding memory leaks.

Initialization of components

Fortran 95 allows us to specify that any object of a derived type is given a default initial value for a component. The value is specified when the component is declared as part of the type definition. If the component is not a pointer, this is done in the usual way. If the component is a pointer, the only initialization allowed is to NULL(). Initialization does not have to apply to all components of a given derived type:

```
TYPE node
   REAL                   :: value = 0.0
   INTEGER                   index
   TYPE(node), POINTER :: next => NULL()
END TYPE node
```

Given an array declaration such as

```
TYPE(node), DIMENSION(100) :: arr
```

sub-objects such as arr(3)%value will have the initial value 0.0, and the value false will be returned by the reference ASSOCIATED(arr(3)%next).

The implications for nested definitions, the interaction with the SAVE attribute, overriding default values, and dynamic object creation are all described in the references above.

Intrinsic procedures

For consistency with similar functions, the numeric intrinsic functions CEILING and FLOOR may now take an optional KIND keyword argument.

The transformational array location functions MAXLOC and MINLOC are extended with an optional argument DIM like those for the functions MAXVAL and MINVAL. The functions MAXVAL, MINVAL, PRODUCT, and SUM already have a DIM argument, but for consistency with MAXLOC and MINLOC, are extended to allow their MASK argument as the second positional argument.

There is a new intrinsic subroutine that returns the processor time, CPU_TIME.

Miscellaneous

There are minor changes concerning comments in NAMELIST input, minimal field width editing, and allowing an optional generic specification on the END INTERFACE statement.

Redundant features

There are six additions to Fortran's list of obsolescent features: fixed source form, computed GO TO, CHARACTER* form of character specification, DATA statements amongst executable statements, statement functions, and assumed character length of function results. Also, for the first time, the standard contains a list of five features that were formerly obsolescent but are now deleted. The references above describe what they are and, for both categories, what to use instead.

4. The future

This article has set out the briefly the changes that Fortran 95 brings to the language, with their justifications. But standardization continues, and Fortran 2000, planned for the year 2001 and whose contents are still speculative, will certainly be described at the appropriate time in these pages.

Michael Metcalf (CERN)