In this issue

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, RG2 9AX, UK

Fax: +44 118 986 9450

Telephone: National 0118 949 9000

International . . . . +44 118 949 9000

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

Cover

The front page feature is a new type of Meteogram which shows the EPS control and members forecast distribution for a model run. The ECMWF graphics package MAGICS was used to produce all the output. MAGICS is described in an article on page 8.

Editorial

ECMWF has been developing versions of its model and analysis with higher vertical resolutions, in particular 50 and 60 levels. The first article in this issue (page 2) describes some of the substantial improvements found in stratospheric analyses and forecasts using these higher resolution versions.

The principal graphics library in use at ECMWF for the past fifteen years has been the Meteorological Applications Graphics Integrated Colour System (MAGICS). The article on page 8 gives the last update on the features and facilities in this library.

On page 15 is the annual review of ECMWF's computer system status and plans.

On 21 September 1998 ECMWF's CFS based data handling system was switched off. The article on page 17 reviews the central role it has played over the years.

One aspect of the work for 1999 will be to prepare for the year 2000 date change. The article on page 20 summarises ECMWF's software status regarding the year 2000.

ECMWF is hosting the second international conference on reanalyses (23-27 August 1999), initial information and a first call for papers can be found on page 21.

Finally see http://www.ecmwf.int/conf/edu/eduprog199901.html for further information on the 1999 training course programme.

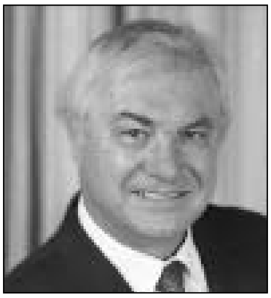

Tony Hollingsworth receives the 1999 Jule G Charney award

For the second time in three years the American Meteorological Society has presented the prestigious Jule G Charney award to an ECMWF staff member.

You will recall that in 1997 Tim Palmer, Head of the ECMWF's Predictability and Diagnostics section received the award *"for fundamental contributions to the theory and practice of extended weather prediction and the understanding of climate phenomena"*

At the January 1999 meeting of the Society Tony Hollingsworth, Head of Research and Deputy Director of ECMWF, received the Charney award for *"penetrating research on four-dimensional data assimilation systems and numerical models"*.

The award is richly deserved and I am delighted that the AMS has chosen to honour Tony's achievements in this manner.

David Burridge

Increased stratospheric resolution in the ECMWF forecasting system

The Centre has been developing versions of its model and data assimilation system with finer and more extensive vertical resolution in both the stratosphere and the planetary boundary layer. In particular, 50- and 60-level versions (differing primarily in their boundary-layer resolution) are being extensively tested in data-assimilation, forecasting and multi-year simulations. The 50-level version is currently undergoing near-real-time trials with a view to early operational implementation, and replacement by the 60-level version is expected later in the year. Indications of some of the substantial improvements found in stratospheric analyses and forecasts are given here. We briefly describe a simple new parametrization of upper-stratospheric moistening by methane oxidation, and discuss the overall simulation of stratospheric humidity, which is strongly dependent on the accuracy of the thermal and dynamical representation of the stratosphere.

The revised versions of the forecasting system

The ECMWF model uses a hybrid vertical coordinate that reduces smoothly from a terrain-following coordinate in the lower troposphere to a pressure coordinate in the stratosphere (Simmons and Burridge, 1981; Simmons and Strüfing, 1983). Since September 1991, a 31-level resolution has been used operationally (Ritchie et al., 1995), with levels distributed as shown in the left-hand portion of Fig. 1. The top four levels are located at pressures of exactly 10, 30, 50 and 70 hPa, and the pressures at the next two levels are very close to 90 and 110 hPa.

The 50-level version of the model is illustrated in the right-hand portion of Fig. 1. The distribution of levels is the same as in the 31-level version below 150 hPa, and levels between 60 and 5 hPa are close to equally-distributed in height with a spacing of 1.5 km. The spacing increases above the 5 hPa level and the top level is at 0.1 hPa.

While this 50-level version of the model was being developed, an independent study of increased vertical resolution in the planetary boundary layer and lower troposphere was yielding promising results (Teixeira, 1999). This led to construction of the 60-level version of the model, with stratospheric resolution similar to that of the 50-level version (but with minor rearrangement of levels), and with much finer resolution in the planetary boundary layer and immediately above. This will form the basis for further developments, as summarized in the concluding section of this article.

A number of problems had to be addressed in developing the versions of the forecasting system with improved stratospheric resolution, and these versions have benefited also from some other recent developments of the forecasting system. Several of the beneficial changes have already been introduced operationally. These include revisions of the radiative parametrization, the two-time-level semi-Lagrangian advection scheme, the calculation of saturation specific humidity at low temperatures, the analysis of humidity and the specification of the ozone climatology.

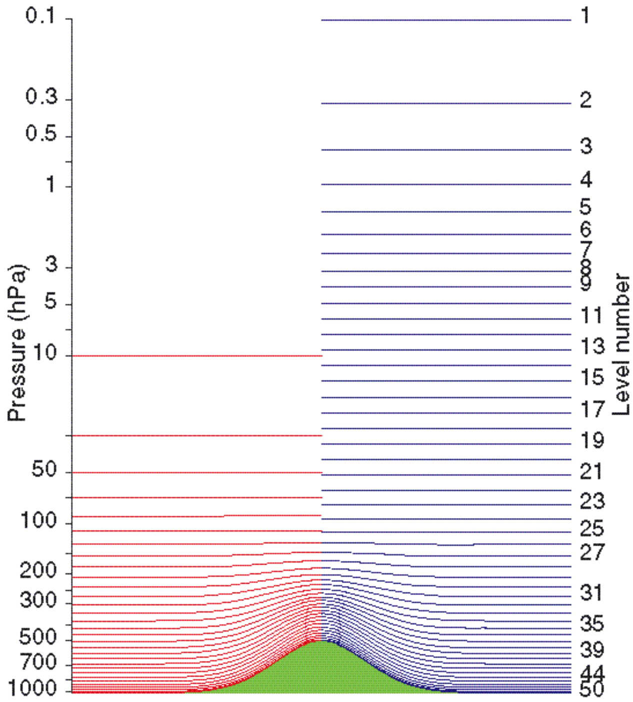

Figure 1: *The distribution of the full model levels at which wind, temperature and humidity are represented, for 31-level (left) and 50-level (right) vertical resolutions, plotted for surface pressures which vary from 1013.25 to 500 hPa.*

Several other changes have been included in the experimental 50- and 60-level systems. Rayleigh friction has been introduced at the uppermost few model levels to ensure a broadly realistic simulation of the mean circulation close to the stratopause. Changes have been made to parameters in the semi-implicit scheme to remove noise arising from weak computational instabilities. Nonlinear normal-mode initialization has been suppressed in the incremental 4D-Var data assimilation (Rabier et al., 1999) to avoid large initialization changes at high levels arising from large normal-mode amplitudes at low pressures. In addition, the model has been enhanced by introduction of the parametrization of methane oxidation specified later in this article.

It was also necessary to compute new sets of balance operators and background error covariances for the data assimilation (Bouttier et al., 1997). These are derived in general from sets of differences between two- and one-day forecasts verifying at the same time. The first set of 50-level forecasts used for these calculations was made from initial conditions formed by merging 31-level ECMWF analyses with UKMO stratospheric analyses (Swinbank and O'Neill, 1994). Implied upper-level errors were reduced to counter effects of incompatibilities between the

ECMWF model and the UKMO analyses. These background statistics were used for two periods of data assimilation and forecasts. Arevised set of statistics was then computed from these 50-level forecasts.

General performance of 50-level system

Data assimilation has been carried out continuously since 15 May 1998 with one form or other of the 50-level system, and 50-level assimilations have also been run for December 1997 and most of January 1998. Objective verification shows that the enhanced resolution provides substantially better stratospheric analyses and forecasts at the levels up to 10 hPa where comparison can be made with results from the standard 31-level system. Fits of 30 hPa temperatures and winds to northern-hemisphere radiosonde measurements are presented in Fig. 2. They are averaged over 165 days of experimentation run with the original background statistics; 50-level experiments for shorter periods using the revised background statistics give slightly better results still. The 50-level analyses (day 0) and forecasts (throughout the range to day 10) can be clearly seen in Fig. 2 to be closer to the verifying radiosonde data than the 31-level analyses and forecasts. Similar plots for other areas and stratospheric levels, and for verification against analyses, confirm the superior performance of the 50-level system. The only notable exception is in the tropical stratosphere, where the 50-level forecasts exhibit a larger growth of bias in temperature, although temperature analyses (and wind analyses and forecasts) are nevertheless better from the 50-level system.

Synoptic study of wintertime stratospheric forecasts also clearly demonstrates that the 50-level system performs better than the 31-level system. The upper panel of Fig. 3 shows the analysis at 10 hPa from the 50-level system for a day on which there was a marked elongation of the polar vortex. The corresponding analysis from the 31-level system (not shown) is very similar. The seven-day forecast verifying on this day from the 50-level system (middle panel) captured quite accurately the elongation of the vortex, whereas the 31-level system (lower panel) erroneously predicted a strong vortex centred over northern Canada and a weaker circulation over Siberia. This is a pronounced case, but the stratospheric forecasts from the 50-level version (or indeed from the 60-level version) are found almost invariably to be better synoptically than those from the 31-level version in depicting features such as the strength and position of the Aleutian high or the shape and orientation of the polar vortex. Dramatic synoptic improvements during a period of major early-winter polar warming have been seen recently in the near-real-time trial of the 50-level system.

Fig. 4 shows the zonal-mean zonal flow for October 1998 from the 50-level analyses and from the UKMO stratospheric analyses. Tropospheric differences should be disregarded, as the averaging was performed on model coordinate surfaces (terrain-following in the troposphere) for the ECMWF analyses and on pressure surfaces for the UKMO analyses. There is a reassuring agreement

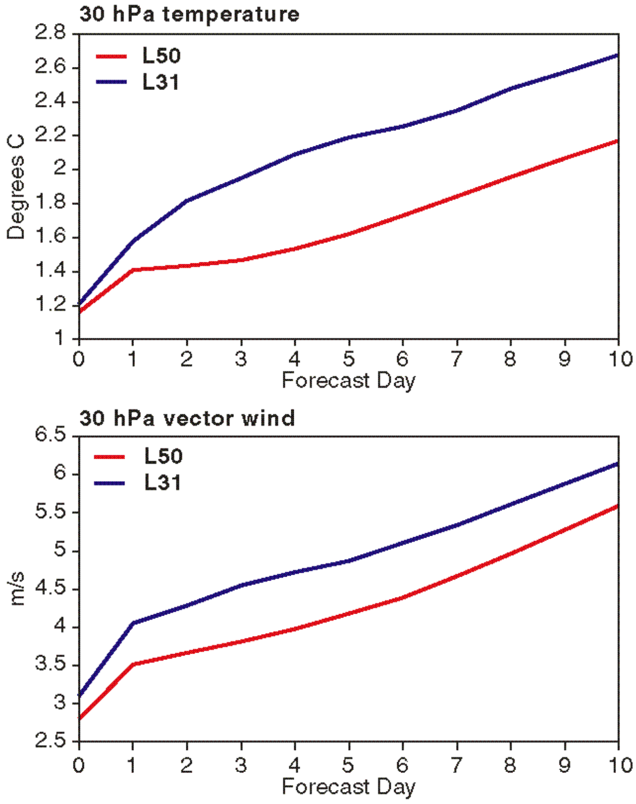

Figure 2: *Root-mean-square errors of 30 hPa temperature and vector-wind analyses (day 0) and forecasts (days 1 to 10) verified against radiosonde measurements over the extratropical northern hemisphere, averaged over a set of 165 cases run with 50-level (red) and 31-level (blue) vertical resolutions.*

between the two mean stratospheric analyses in the extratropics, where the only difference of note occurs close to the stratopause in the southern hemisphere. Differences are, however, more marked in the tropics. The easterly maximum located at about 40 hPa is more than 5 ms^{-1} stronger in the ECMWF analysis, and this analysis has westerlies of up to 25 ms^{-1} above 10 hPa, where the UKMO analysis has easterlies. It is not surprising that differences are large in this region, because of difficulty in modelling the quasi-biennial oscillation (QBO) because of the paucity of wind observations, and because of the limited extent to which the wind analyses in the tropics can be controlled by assimilation of satellite radiance measurements. There is evidently a need for further study of the realism of the tropical stratospheric wind analyses from the new system, particularly to examine performance over a complete cycle of the QBO.

Figure 5 presents objective verification of 500 hPa height forecasts from the 50- and 31-level systems. It is based on the extensive sets of forecasts available with the original background statistics; the average is over the 154 forecasts that can be verified against analyses produced using the same model version that provided the initial conditions. The 50-level version gives better mean verification scores throughout the medium range.

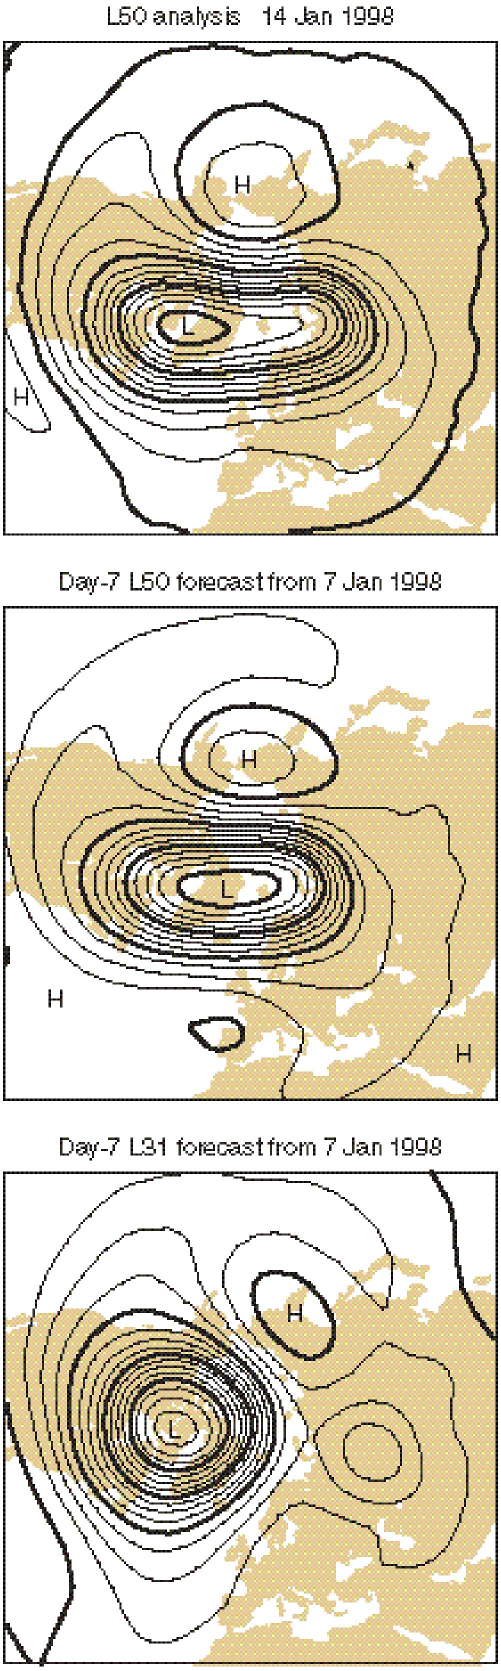

Figure 3: *The 10 hPa height analysis for 12UTC 14 January 1998 from the 50-level version (upper), and seven-day forecasts verifying on this date from the 50-level (middle) and 31-level (lower) versions of the forecasting system. The contour interval is 160 m.*

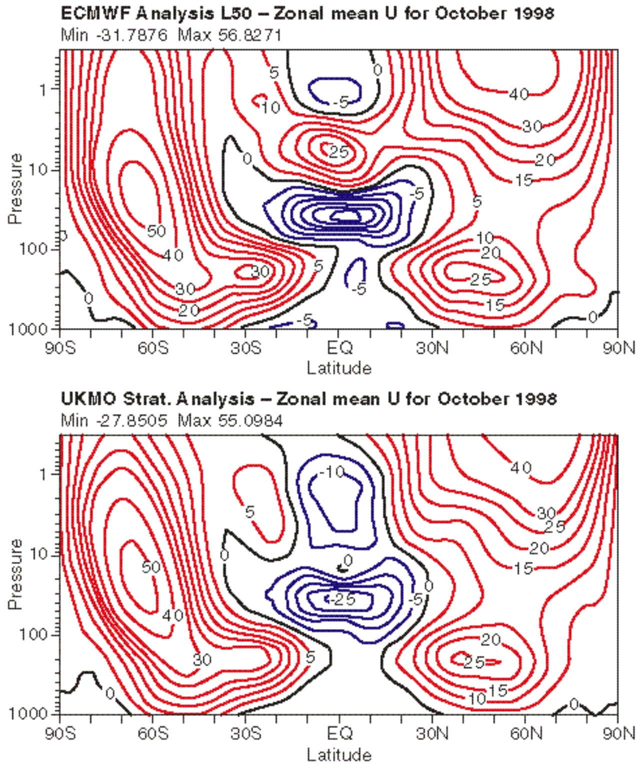

Figure 4: *Meridional cross-sections of zonal-mean zonal wind averaged for the month of October 1998 from experimental 50-level ECMWF analyses (upper) and from the stratospheric analysis system of UKMO (lower). Red contours denote westerlies, blue contours easterlies, and the contour interval is 5 ms^{-1}. Tropospheric differences should be disregarded.*

Although differences appear small, they are in fact quite significant when interpreted in the context of the long-term evolution of forecast skill, which has seen an increase by the order of one day per ten years in the forecast range at which the anomaly correlation falls below 60%.

Repeating a subset of the forecasts included in Figure 5 using the revised background statistics yields very similar hemispheric verification scores for the 50-level system, but some marked regional differences. In particular, much less improvement was found over Europe with the revised background statistics, whereas the converse was the case for North America. Further study of the specification of the background statistics is evidently needed and is planned (see later). For the time being, the 50-level system with revised background statistics currently undergoing near-real-time parallel trials is expected to yield substantial improvements in operational stratospheric products and some small improvement in tropospheric products, with further improvements expected from additional work to be carried out after operational implementation.

The treatment of methane oxidation

ECMWF analyses and simulations have typically been too dry in the stratosphere, because of the absence of a representation of the high-level source of water vapour due to

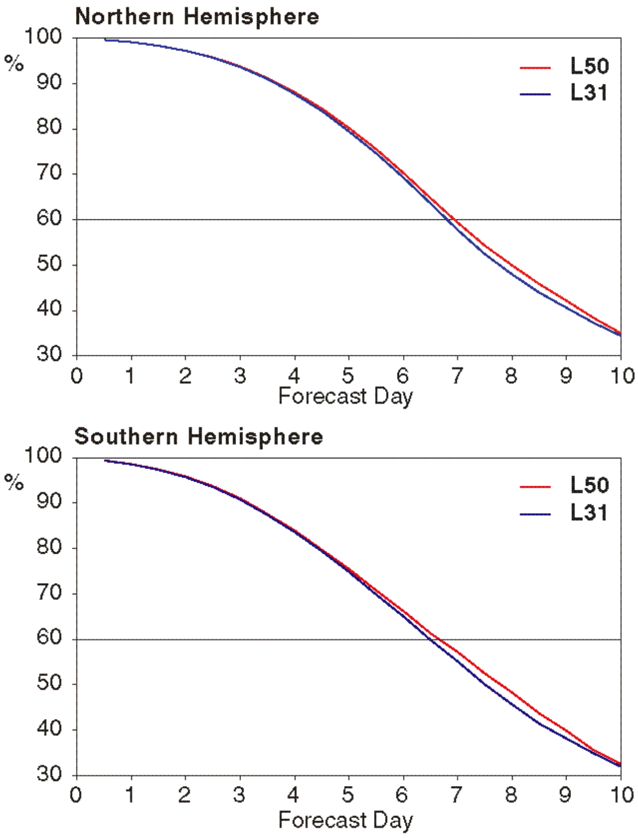

Figure 5: *Anomaly correlations of 500 hPa height averaged over 154 cases run with 50-level (red) and 31-level (blue) vertical resolutions, for the extratropical northern (upper) and southern (lower) hemispheres.*

the oxidation of methane (Simmons et al., 1999). This has been seen, for example, in a recent 19-year simulation using observed surface boundary conditions for the years 1979-1997, carried out as a contribution to the second phase of the Atmospheric Model Intercomparison Project (AMIP; Gates, 1992). Arun of the 50-level version of the model using T63 horizontal resolution dried gradually over the first decade of integration, leading to upper stratospheric specific humidities around 1.7-1.8 mg/kg. These values were maintained over the second half of the integration period. This is not unreasonable in the absence of methane oxidation, as it corresponds to a long-term water-vapour mixing ratio of around 2.8 ppmv, giving a value of 6.2 ppmv when twice the tropospheric methane mixing ratio of 1.7 ppmv is added. Photochemical calculations and observations (e.g. Brasseur and Solomon, 1984; Bithell et al., 1994) indicate that the sum of the mixing ratio of water vapour and twice that of methane is approximately constant over much of the stratosphere, with a value close to or a little above 6 ppmv.

Asimple parametrization of the moistening by methane oxidation has thus been developed. The basic assumptions are that the volume mixing ratio of water vapour $[H_2O]$ increases at rate $2k_1[CH_4]$ and that there is a steady balance between the mixing ratios of methane and water vapour:

$$2[CH_4] + [H_2O] = 6\ ppmv$$

The rate of increase of water vapour (in ppmv) is then

$$k_1(6 - [H_2O])$$

In terms of specific humidity, q , the source is

$$k_1(Q-q)$$

where Q has the value 3.75 mg/kg.

For completeness, an extra photolysis term $-k_2q$ is included in the mesosphere, although it has little effect for the 50- and 60-level resolutions discussed here. k_1 and k_2 are specified as functions of pressure, with k_1 equal to (100 days)$^{-1}$ at pressures less than 1 hPa. The vertical profiles of k_1 and k_2 are chosen such that the dependence on altitude of the combined photochemical lifetime, $(k_1 + k_2)^{-1}$, shown in Figure 6, is similar to that presented by Brasseur and Solomon (1984). The slow time scale of the process in the stratosphere enables latitudinal and temporal variations in relaxation rate to be neglected in the first instance.

Stratospheric humidity in simulations and analyses

A new AMIP simulation including the parametrization of methane oxidation has been run using the 60-level vertical resolution. After an initial adjustment period of several years, the stratospheric humidity in this simulation exhibits a fairly regular annual cycle. Latitude/pressure sections showing 15-year means for January, April, July and October are presented in Figure 7. These sections may be compared with observational data from the UARS satellite processed by Randel et al. (1998). The model successfully captures the dryness both of air entering the stratosphere in the tropics in the boreal winter and of air in the cold Antarctic lower stratosphere in austral winter and spring, this being consistent with realistic simulated temperatures at the tropical tropopause and in the polar winter stratosphere. The consequences of a general ascent of relatively dry air throughout the tropical stratosphere, and the net high-latitude wintertime descent (and summertime ascent) of air moistened by methane oxidation can be clearly seen. The model is unable to capture the observed dryness of the winter polar upper stratosphere as it does not have the resolution to describe correctly the wintertime descent of dry air from the mesosphere. It represents the upward tropical transfer of the annual

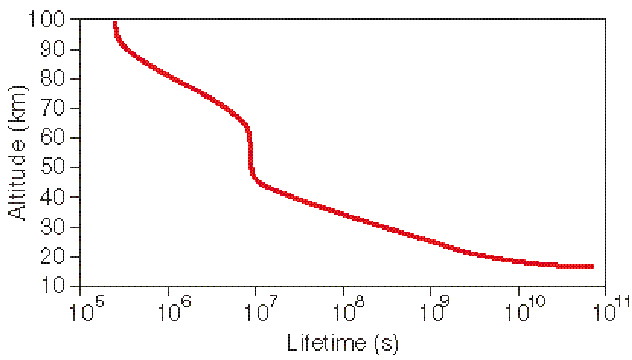

Figure 6: *Combined photochemical lifetime.*

cycle in water vapour (the "tape-recorder" effect discussed by Mote et al., 1996), but is generally too diffusive, and drier in the tropical middle stratosphere by some 10-20% compared with UARS data.

No observations of stratospheric humidity are used in the ECMWF analysis system. Humidity simply evolves during the assimilation according to the model's dynamical and parametrized physical processes, with winds and temperatures (and tropospheric humidity) constrained by the analysis process. The parametrization of methane oxidation was activated on 17 June 1998 in the data assimilation cycles using the 50-level version of the model. The stratospheric analyses have not yet become fully adapted to the change, but it is nevertheless instructive to compare 31-level and 50-level analyses. Meridional cross-sections of zonal-mean specific humidity averaged over the last week of November and the first week of December are shown in Figure 8.

The parametrization of methane oxidation has had little effect below 10 hPa by the time shown in Figure 8. Lower stratospheric differences instead arise mostly from the resolution difference. The 31-level analyses generally exhibit a too rapid upward transfer of relatively moist or dry air introduced at the tropical tropopause (Simmons et al., 1999). 50-level simulations are found to be much more realistic in this respect, although upward transfer and attenuation in the deep tropics are still stronger than observed, as in the 60-level AMIP simulation discussed above. Figure 8 shows that the benefit seen in the high-resolution simulations carries over into the 50-level analyses. Relatively moist air was introduced into the stratospheric analyses in the boreal subtropics over the summer of 1998, and by December the maximum in humidity has reached 10 hPa in the 31-level analyses, but is still below 30 hPa in the 50-level analyses. The maximum is located further north in the 31-level system, and much more moistening of the northern extratropics has taken place in this system. At high southern latitudes, the 31-level system is moister than the 50-level system in the lower stratosphere, but drier at 10 hPa. The latter is due partly to colder 10 hPa temperatures (and more condensation) in winter and early spring in the 31-level system and partly to descent of moister air from the upper stratosphere in the 50-level system. By December, however, there is mean ascent at high southern latitudes, this acting to dry the middle stratosphere in this region and season.

Daily analyses of humidity over the southern hemisphere from 30 November to 5 December 1998 are presented in Figure 9, for the model level close to 51 hPa from the 50-level system. The driest air is located in the decaying, perturbed westerly vortex, where dehydration by the parametrized condensation process occurred earlier when temperatures were sufficiently cold. The maps show evolving filamentary structure in the humidity field, and indications of mixing. Observations of fields which behave largely as passive tracers can be used effectively to extract information on

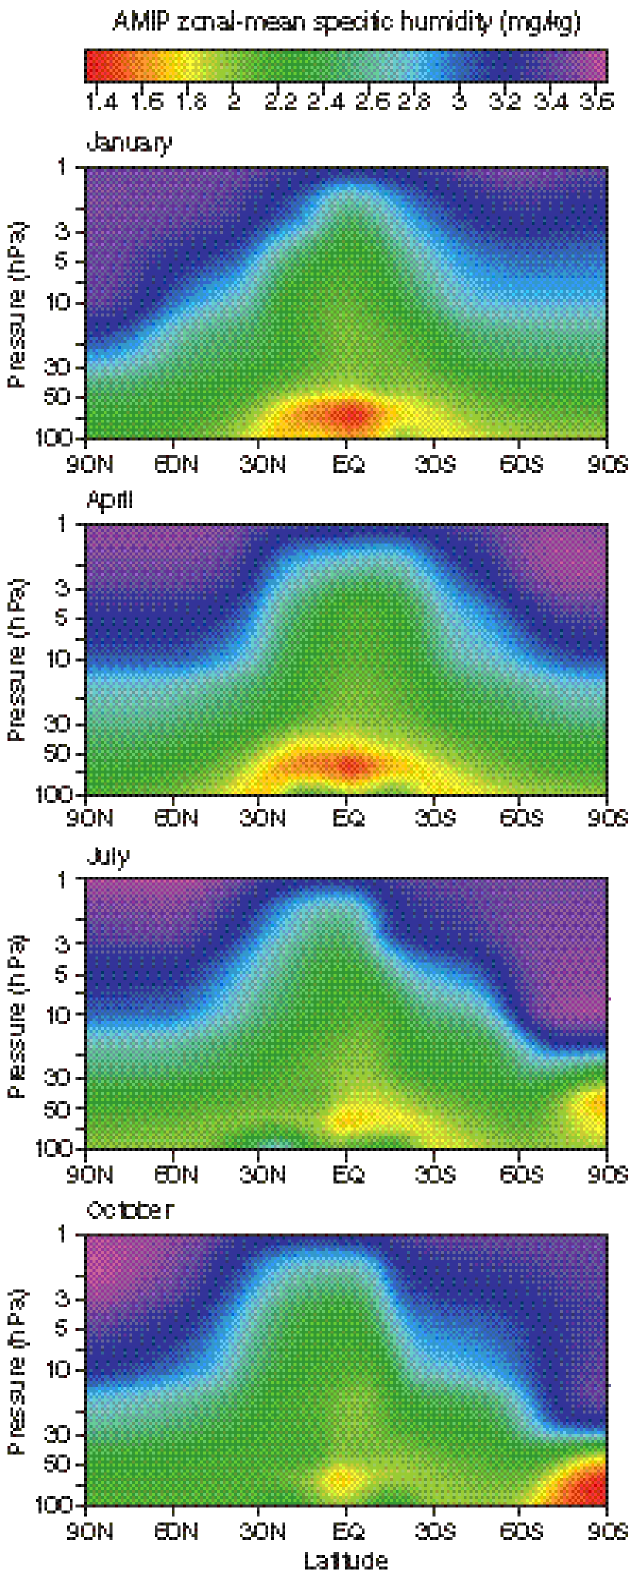

Figure 7: *Meridional cross-sections of zonal-mean specific humidity for January, April, July and October from a 60-level AMIP simulation. Each section is a 15-year average for 1983-1997.*

the wind field in a 4D-Var assimilation system. This is a prime motivation for ongoing work to include stratospheric ozone as a variable of the ECMWF system.

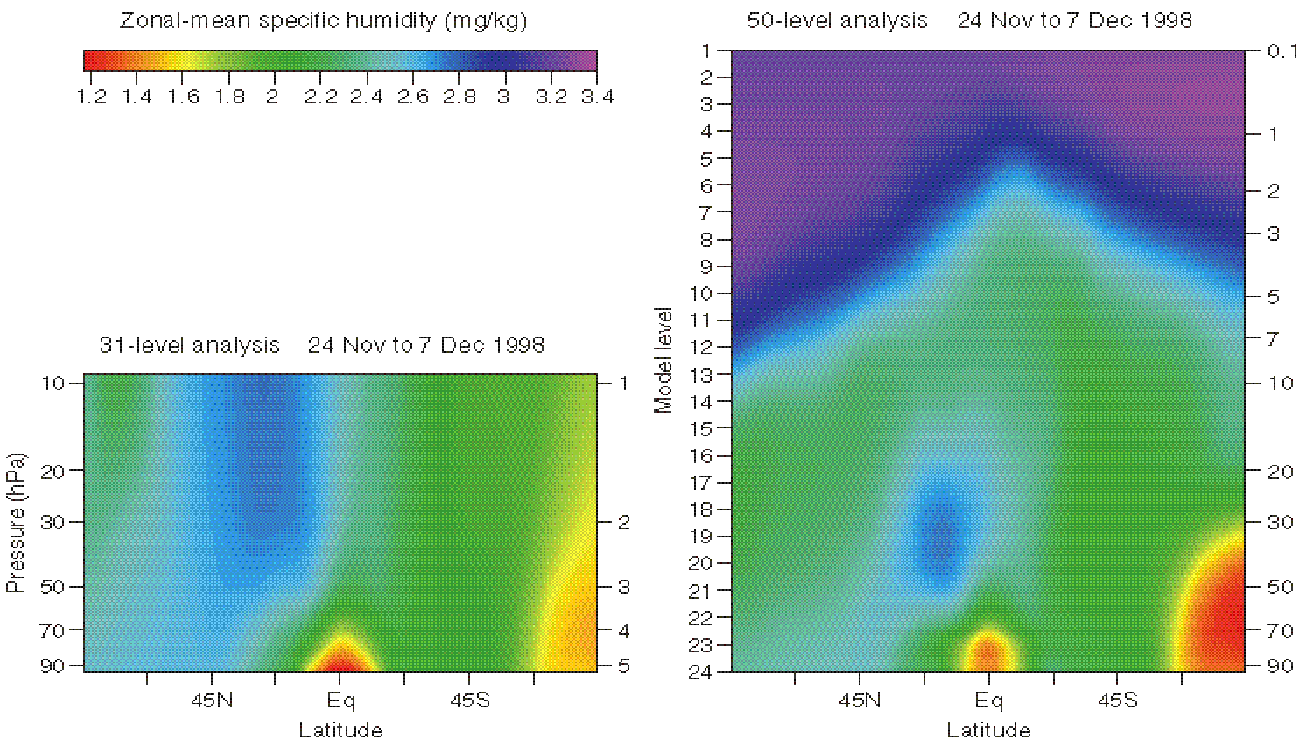

Figure 8: *Meridional cross-sections of zonal- and time-mean specific humidity from 31- and 50-level analyses.*

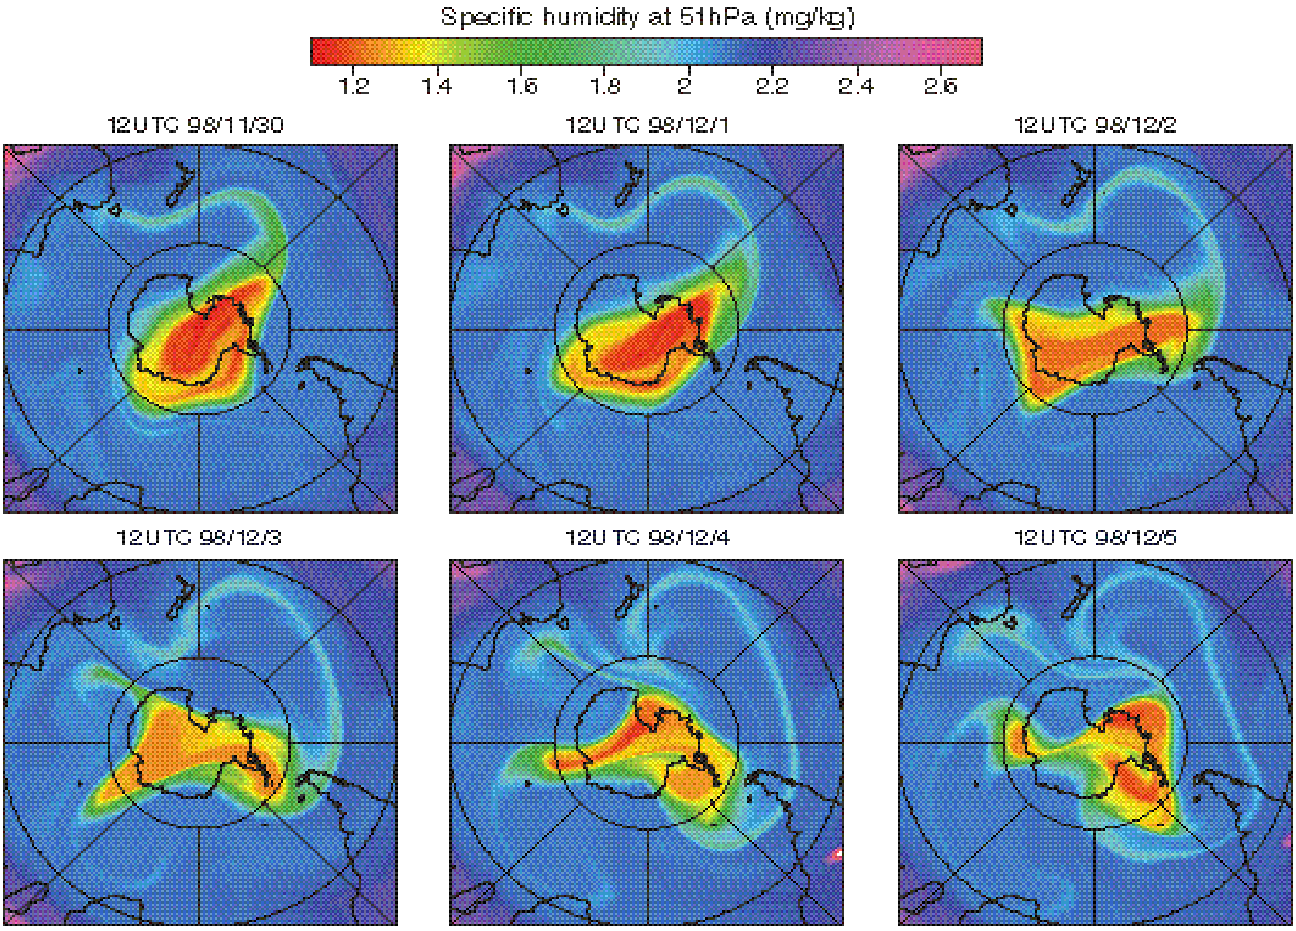

Figure 9: *Specific humidity at close to 51 hPa from 50-level analyses for consecutive days from 30 November to 5 December 1998.*

Plans

Our immediate goal is the operational implementation of the 50-level version of the forecasting system. Included with this will be a new fast radiative transfer model for the assimilation of radiance data from satellites (Saunders et al., 1999). The availability of model background fields that cover the whole of the stratosphere as well as the troposphere will be exploited in a new scheme for the assimilation of raw radiances from the TOVS and new ATOVS instruments. It is expected that this scheme will be introduced into operations soon after the 50-level system.

Further development is being concentrated on the 60-level system. The specification of background error statistics will be investigated in conjunction with a revision of observation errors which addresses known deficiencies. Assimilation of ozone data will be added to the 60-level system. Operational implementation is planned for later this year. Work is also in hand to use the 60-level system for the new ECMWF reanalysis, ERA-40, covering the period from 1957 to the present day.

Acknowledgements

This work has been greatly facilitated by the use of stratospheric analyses made available by the UK Meteorological Office. Ideas and advice from Bob Harwood and colleagues at Edinburgh University on the representation of methane oxidation are gratefully acknowledged.

References

Bithell, M., **Gray**, **L.J.**, **Harries**, **J.E.**, **Russell**, **J.M. III**, and **Tuck**, **A.F.** (1994) Synoptic interpretation of measurements from HALOE. *J. Atmos. Sci.*, **51**, 2942-2956.

Bouttier, **F.**, **Derber**, **J.**, and **Fisher**, **M.** (1997) The 1997 revision of the J_b term in 3D/4D-Var. *ECMWF Tech. Memo.*, **238**, 54 pp.

Brasseur, **G.**, and **Solomon**, **S.** (1984) Aeronomy of the Middle Atmosphere. *D. Reidel Publishing Co.*, Dordrecht,452 pp.

Gates, **W.L.** (1992) AMIP: The Atmospheric Model Intercomparison Project. *Bull. Amer. Meteorol. Soc.*, **73**, 1962-1970.

Mote, P.W., **Rosenlof**, **K.H.**, **McIntyre**, **M.E.**, **Carr**, **E.S.**, **Gille**, **J.C.**, **Holton**, **J.R.**, **Kinnersley**, **J.S.**, **Pumphrey**, **H.C.**, **Russell**, **J.M. III**, and **Waters**, **J.W.** (1996) An atmospheric tape recorder: The imprint of tropical tropopause temperatures on stratospheric water vapour. *J. Geophys. Res.*, **101D**, 3989-4006.

Rabier, **F.**, **Järvinen**, **H.**, **Klinker**, **E.**, **Mahfouf**, and **Simmons**, **A.** (1999) The ECMWF operational implementation of four dimensional variational assimilation. Part I: Experimental results with simplified physics. Submitted to *Q. J. R. Meteorol. Soc.*.

Randel, **W.J.**, **Wu**, **F.**, **Russell**, **J.M. III**, **Roche**, **A.**, and **Waters**, **J.W.** (1998) Seasonal cycles and QBO variations in stratospheric CH_4 and H_2O observed in UARS HALOE data. *J. Atmos. Sci.*, **55**, 163-185.

Ritchie, **H.**, **Temperton**, **C.**, **Simmons**, **A.**, **Hortal**, **M.**, **Davies**, **T.**, **Dent**, **D.**, and **Hamrud**, **M.** (1995) Implementation of the semi-Lagrangian method in a high resolution version of the ECMWF forecast model. *Mon. Wea. Rev.*, **123**, 489-514.

Saunders, **R.W.**, **Matricardi**, **M.**, and **Brunel**, **P.** (1999) An improved fast radiative transfer model for assimilation of satellite radiance observations. *Q. J. R. Meteorol. Soc.*, **125**, in press.

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

Simmons, **A.J.**, and **Strüfing**, **R.** (1983) Numerical forecasts of stratospheric warming events using a model with a hybrid vertical coordinate. *Q. J. R. Meteorol. Soc.*, **109**, 81-111.

Simmons, **A.J.**, **Untch**, **A.**, **Jakob**, **C.**, **Kållberg**, **P.**, and **Undén**, **P.** (1999) Stratospheric water vapour and tropical tropopause temperatures in ECMWF analyses and multi-year simulations. *Q. J. R. Meteorol. Soc.*, **125**, 353-386.

Swinbank, **R.**, and **O'Neill**, **A.** (1994) A stratosphere-troposphere data assimilation system. *Mon. Wea. Rev.*, **122**, 686-702.

Teixeira, **J.** (1999) The impact of increased boundary layer vertical resolution on the ECMWF forecast system. *ECMWF Tech. Memo.* 267.

Agathe Untch, Adrian Simmons and colleagues

MAGICS - The ECMWF graphics package

The Meteorological Applications Graphics Integrated Colour System (MAGICS), developed at the European Centre for Medium-Range Weather Forecasts (ECMWF) is a software system for plotting contours, satellite images, wind fields, observations, symbols, streamlines, isotachs, axes, graphs, text and legends.

Development started in 1984 with the first release for ECMWF staff in 1985. It has since been used by most Member States. It is also used in the meteorological workstation systems Metview[1] and Synergie[2].

MAGICS was designed to conform to new meteorological and graphics standards, e.g. GRIB[3], BUFR[4]. The ECMWF software package EMOSLIB is used to handle the processing of GRIB and BUFR data. An advanced contouring method (CONICON[5]) is used.

Data to be plotted may be presented in various formats, i.e. GRIB code, BUFR code or in matrices.

The current version of MAGICS, version 6, has been developed with three device drivers. The MAGICS/PostScript version produces PostScript output, e.g. in

1 See ECMWF Newsletter number 68 winter 1994/95

2 Developed by Météo France

3,4 WMO standard

5 Developed by the University of Bath, UK

batch, and does not assume any underlying graphics package. MAGICS/OpenGL and MAGICS/METVIEW drivers are aimed at Metview usage.

MAGICS has been fully tested for Year 2000 compliance.

The MAGICS Reference Manual contains full documentation on the use of MAGICS. Further documentation is available in the MAGICS WWW pages.

This article is an update of the article published in the Winter 1995/96 issue of the ECMWF Newsletter.

MAGICS Features

MAGICS contains the following features:

- Flexible user interface with a comprehensive list of simple English language parameters;
- Extensive use of colour;
- Matrix, GRIB and BUFR code data input;
- Selection of geographical area and direct projection of data;
- Polar stereographic, cylindrical, Mercator and satellite projections;
- High quality contouring based on CONICON and lower quality linear contouring;
- Satellite image plotting in all four projections;
- Shading between contour lines and also in graph and coastline plotting;
- Wind fields directly projected, as arrows, WMO flags, streamlines or isotachs;

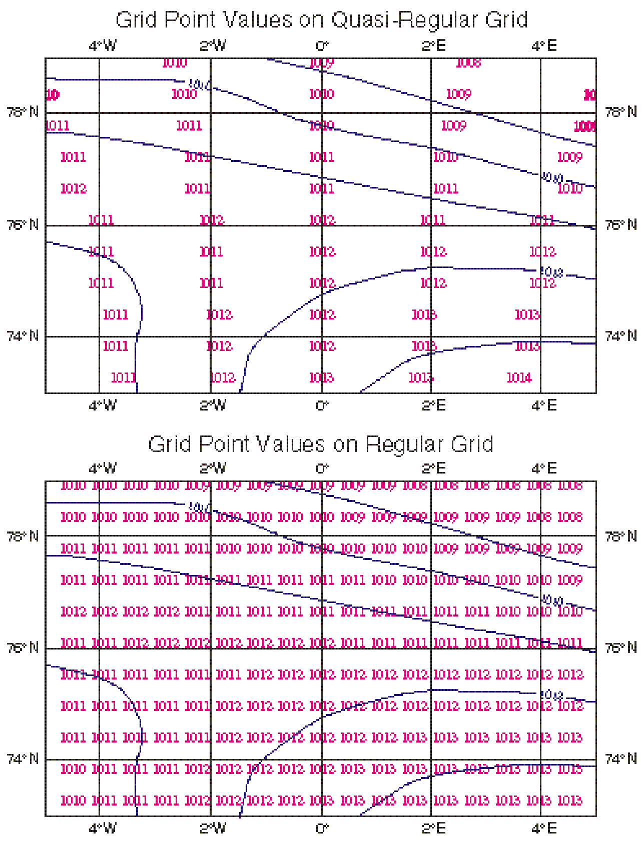

***Figure 1:** Plots of grid point values, one from the original quasi-regular Gaussian grid and one from the same field interpolated onto a regular Gaussian grid.*

- Full observation plotting;
- Symbol plotting;
- Axis and graph plotting;
- Legend plotting;
- Storage of program context (specification groups).

MAGICS subroutines and parameters

MAGICS consists of a small set of FORTRAN callable subroutines which are divided into five different categories: initialisation, parameter setting, enquiry, action and pseudo action.

Parameters in MAGICS are the attributes to be assigned to the various items that make up plotted output, such as line-style, colour and size of plot. They control all plotted output. The parameters may be set by calling MAGICS parameter setting routines, where the use of a keyword (a MAGICS parameter) as an argument defines the action to be taken. Action routines are the only ones that produce plotted output and hence the parameters associated with each action routine should be set, if required, before the action routine is called.

A typical MAGICS program, as in the simple example below, consists of a series of parameter setting, action and pseudo action routines.

```
CALL POPEN                                    open the MAGICS package
CALL PSETC ('PS_DEVICE', 'ps_vax_c')                  set output device
CALL PSETC ('GRIB_INPUT_FILE_NAME', 'data.grb')  pass the GRIB data for contouring
CALL PGRIB                                          decode the GRIB data
CALL PSETC ('CONTOUR_LINE_COLOUR', 'RED')            parameter setting
CALL PCOAST                            default map is global cylindrical
CALL PCONT                                               plot contours
CALL PCLOSE                                               close MAGICS
```

Plot layout

To make the positioning of users' plots simple, MAGICS has introduced the concepts of subpage, page and super page.

A super page corresponds to the plotting area to be used on a plotting device. With the plot layout facilities, pages can be positioned within the super page either automatically or by the user. Similarly, one or more subpages can be plotted inside a page.

Mapping

Mapping is the placing or projecting of coastlines, grids and data onto the user's subpage. Four types of projections are catered for in MAGICS, cylindrical, polar stereographic, Mercator and satellite. It is not necessary for users to extract the required area before passing data, as MAGICS can perform this function and the conversion to the various projections. It is also possible to shade both the land and sea areas (see Figure 3).

Data input

Data to be plotted can be presented in a number of ways, depending on its type: contour fields, wind fields, observations, satellite images, axis and graph data etc.

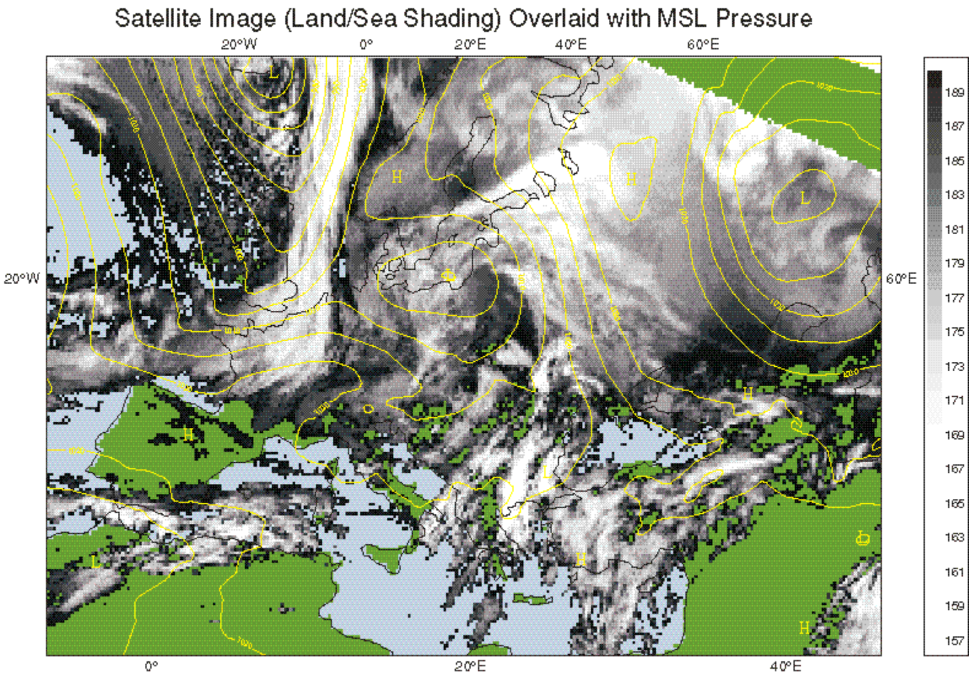

Figure 2: *Meteosat Infra - red satellite image 12th January 1999 1200 UTC overlaid with contours of MSL pressure, ECMWF analysis 12th January 1999 1200 UTC*

Acontour field can be presented as either a two-dimensional array of grid point values or in GRIB code format. GRIB code can be presented as either an external file or as a one-dimensional data array. Matrix data can be regularly spaced or on a Gaussian grid.

A wind field can be presented as either a pair of two-dimensional arrays of grid point values or in GRIB code format. The pair of two-dimensional arrays can contain either U or V velocity components, or speed and direction components. Data in GRIB code can be presented as either an external file or as a pair of one-dimensional data arrays.

Observations must be presented as an external file. Data for observation plotting must be in WMO BUFR code format.

Satellite images must be presented in GRIB code format.

Axis and graph data must be presented in one dimensional arrays.

GRIB data

Facilities exist in MAGICS for the automatic processing and plotting of GRIB coded data. In particular MAGICS can decode GRIB data, convert from spectral to grid point, scale the fields and set up the MAGICS data input parameters. The pseudo action routine PGRIB handles all these functions. PGRIB can handle data that is in either GRIB edition 0 or 1, various Gaussian grids and bit-map fields. It can also cope with pseudo (quasi-regular) GRIB code, extended GRIB code for satellite image data, GRIB data on stretched/rotated grids and GRIB data in polar stereographic projection. Aspecial feature allows the plotting of wave direction and heights as wind arrows.

MAGICS allows users to develop and use an extended version of GRIB Table 2. The extended version can include parameter information, i.e. text, scaling of values, units used, contour interval etc. This enables users to define the local use of weather parameters.

Contouring

MAGICS contouring is primarily based on CONICON, which is designed to draw the contours of continuously differentiable fields (see Figures 1, 2 and 6). There is also a faster but coarser method of contouring available in MAGICS, based on a linear interpolation.

For the user interface, emphasis has been placed on simplicity, making it easy for the user to define the contour levels required and the attributes to draw them with. Users have full control over plotting of labels, highs and lows as well as the thickness and colour of contour lines.

MAGICS shading facilities enable the user to shade between contour levels (See section on MAGICS Shading). Another feature of MAGICS contouring allows users to plot the grid point values on their exact location, either on their own or overlaid on contour lines. For pseudo-GRIB data, the grid point values can be plotted on the reduced grid or on the interpreted grid. Examples of grid point plotting are given in Figure 1.

MAGICS allows the plotting of `split' level contours, where contours above a certain level can be plotted with different characteristics to those contours below the specified level.

Satellite image plotting

MAGICS provides facilities for plotting satellite images (see Figure 2) with optional overlaid fields, observations and coastlines. Each pixel in an image will be plotted in the tile format, i.e. each pixel is represented by a small rectangle filled with a colour (note that grey scale definitions are special cases of colour definitions). Input image data must be in the spaceview projection. If necessary, MAGICS will automatically reproject the image from spaceview to the required projection.

Image data must be passed to MAGICS in the extended GRIB representation.

The colour distribution of the plotted image is controlled by the use of a colour table. The user may specify his own colour table or use the MAGICS default one.

Wind field plotting

Wind fields may be presented to MAGICS as U and V velocity components, as speed and direction components, or as GRIB code data. They may be plotted in one of four ways:

- **wind arrows**: where the wind is represented by a vector whose length is proportional to the speed of the wind;
- **WMO standard wind flags**: where a wind flag is represented by barbs and solid pennants;
- **streamlines**: Astreamline is a line whose tangent at any point is parallel to the instantaneous velocity at that point;
- **isotachs**: which are contour lines of equal wind speed.

Wind related parameters, which are generated from differential properties of the wind field, may be plotted using normal contouring facilities, e.g. Divergence, Vorticity, Stream Functions and Velocity Potential. The use of colour enables the plotting of coloured wind arrows where the colour of the arrow can represent the relevant temperature, humidity, etc.

Figure 3 shows examples of coloured winds, streamlines, wind flags and wind arrows.

Observation plotting

All observation types can be plotted, including sea-wave and swell heights. A special feature of observation plotting is the facility to plot information from the ECMWF Analysis Feedback (see Figure 4), i.e. rejected data, data used or not used by the analysis and departure values. WMO standard plotting formats are adhered to as closely as possible.

The user has full control over the type, number and size of the observations to be plotted. The amount of meteorological information plotted may also be controlled by the user, for example it is possible to plot only the positions of observations.

MAGICS allows the use of BUFR parameter tables which enable the user to locally define parameter descriptor numbers.

Axis plotting

MAGICS axis facilities allow users to plot vertical and horizontal axes. These facilities include axis labelling, axis title plotting and subdivision of axes with ticks. MAGICS Axis parameters give the user control over all

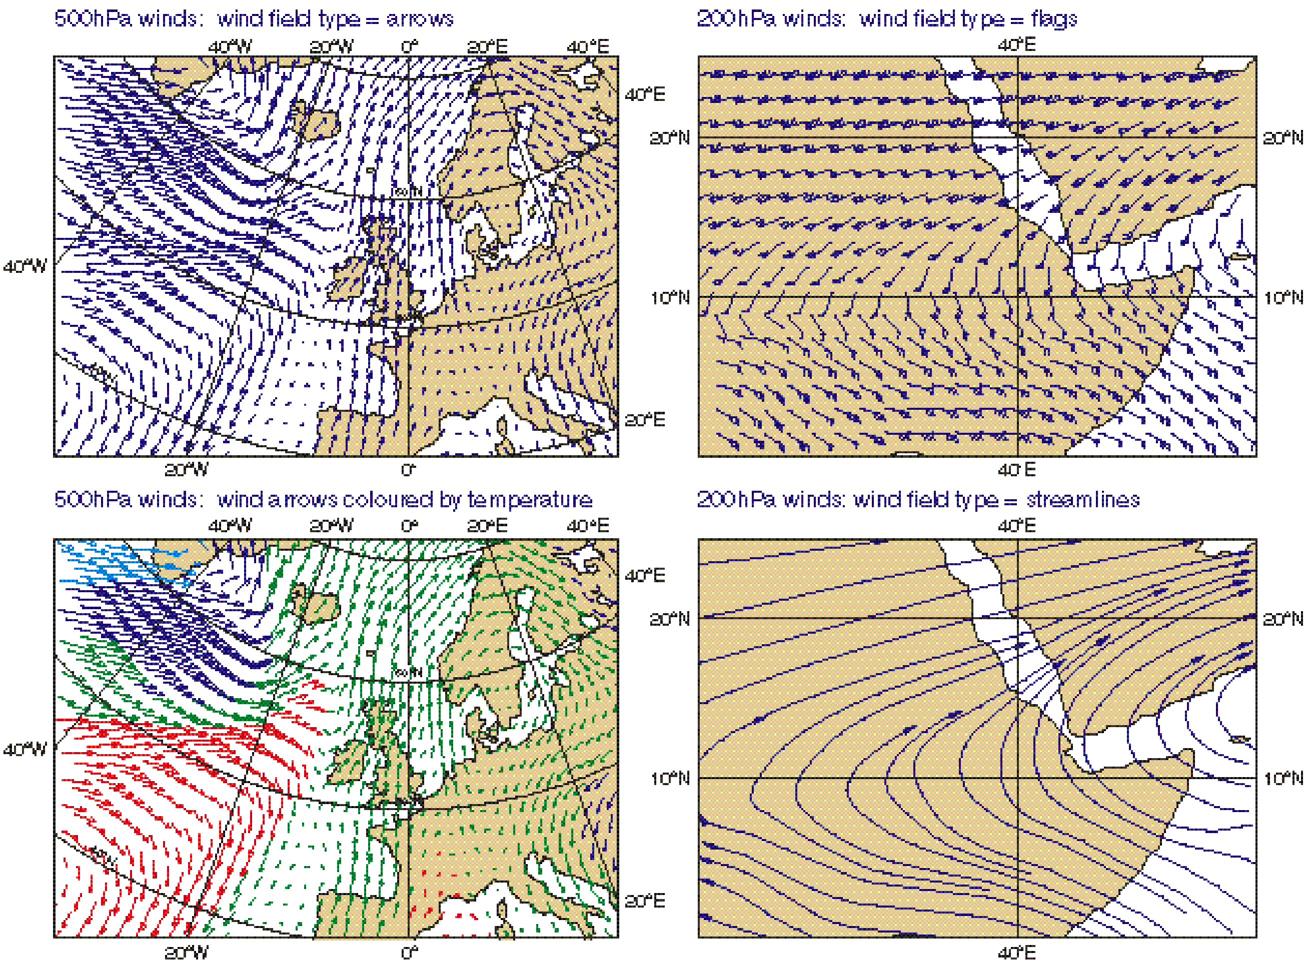

Figure 3: *Examples of wind representations in MAGICS.*

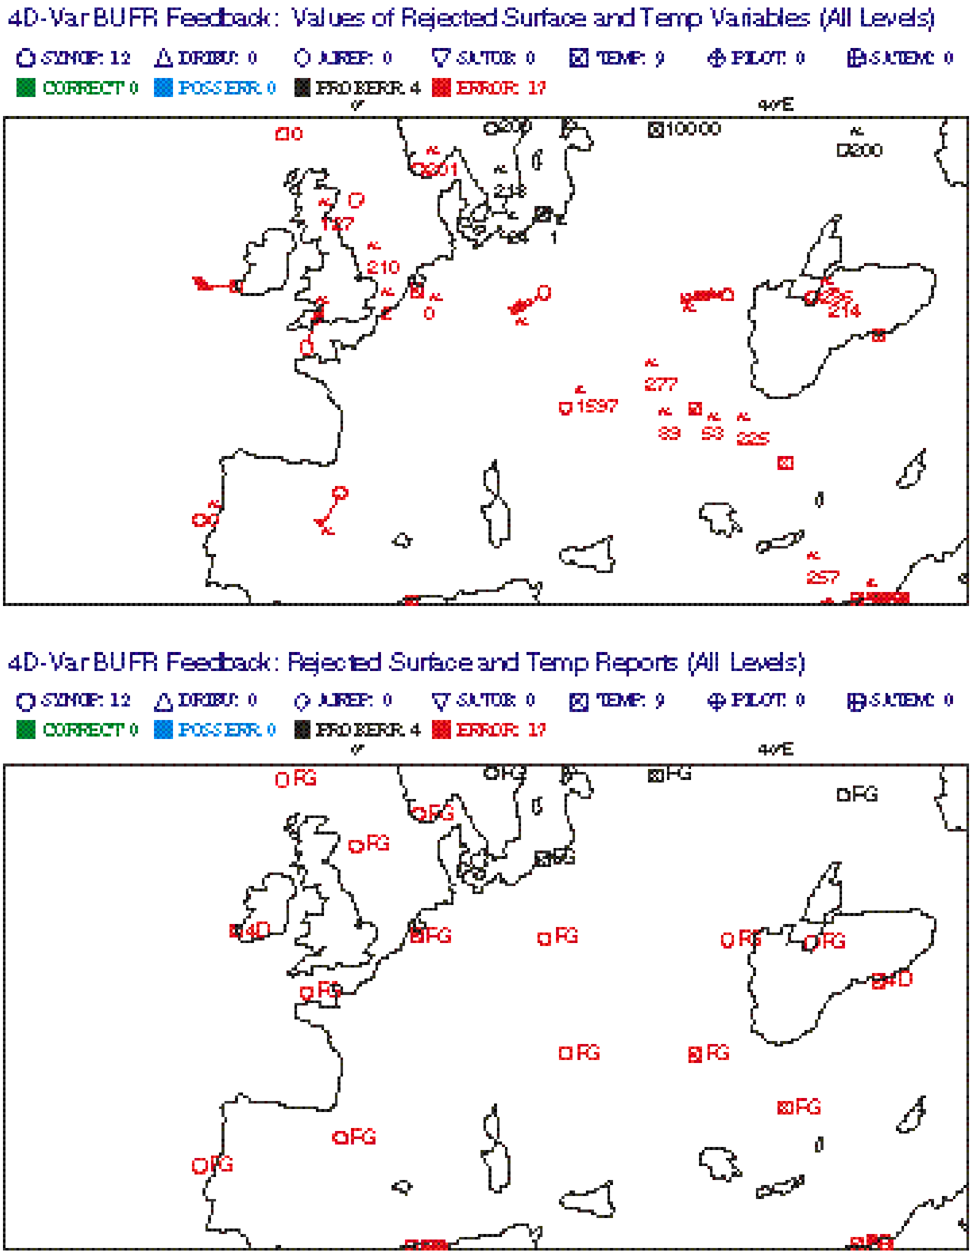

***Figure 4:** Plots of 4D-Var Feedback information from ECMWF Analysis associated with observations over Europe.*

aspects of axis plotting, e.g. position, orientation, colour, thickness, etc. Axes may be regular, logarithmic or user-defined and may be subdivided with tick marks. Further subdivision between ticks, i.e. minor ticks, is also possible. Ticks may be labelled either on the tick mark or between ticks. Grid lines may be plotted. MAGICS axis facilities can be used for plotting cross-sections or for graph plotting.

A special feature in Axis plotting is the automatic DATE/TIME axis facility (see front cover). If the user specifies the start and end date, MAGICS will generate the axis with the correct hour, day, month and year.

Graph plotting

Graph plotting in MAGICS is the plotting of line charts (curves), bar charts and area charts within a set of axes (see front cover). An area chart is where the area between two curves is shaded. The user has full control over line style, colour and thickness of graphs and, if required, a legend describing the graph can be plotted. Bar charts and area charts can use all the MAGICS shading facilities.

Symbols can be plotted on the curves and a specified area around curves can be blanked. Curves can be drawn straight or rounded, where a smoothing algorithm is applied to the curve. There are options for dealing with missing data.

MAGICS can automatically generate the axes for GRAPH plotting by analysing the input data. Another feature is the facility for drawing curves on maps.

Symbol plotting

Symbol plotting in MAGICS is the plotting of different types of symbols at selected locations (see Figure 5). A symbol in this context is a number, a text string, a WMO wind flag or a MAGICS marker. The position of a symbol may be defined by its geographical location (latitude/

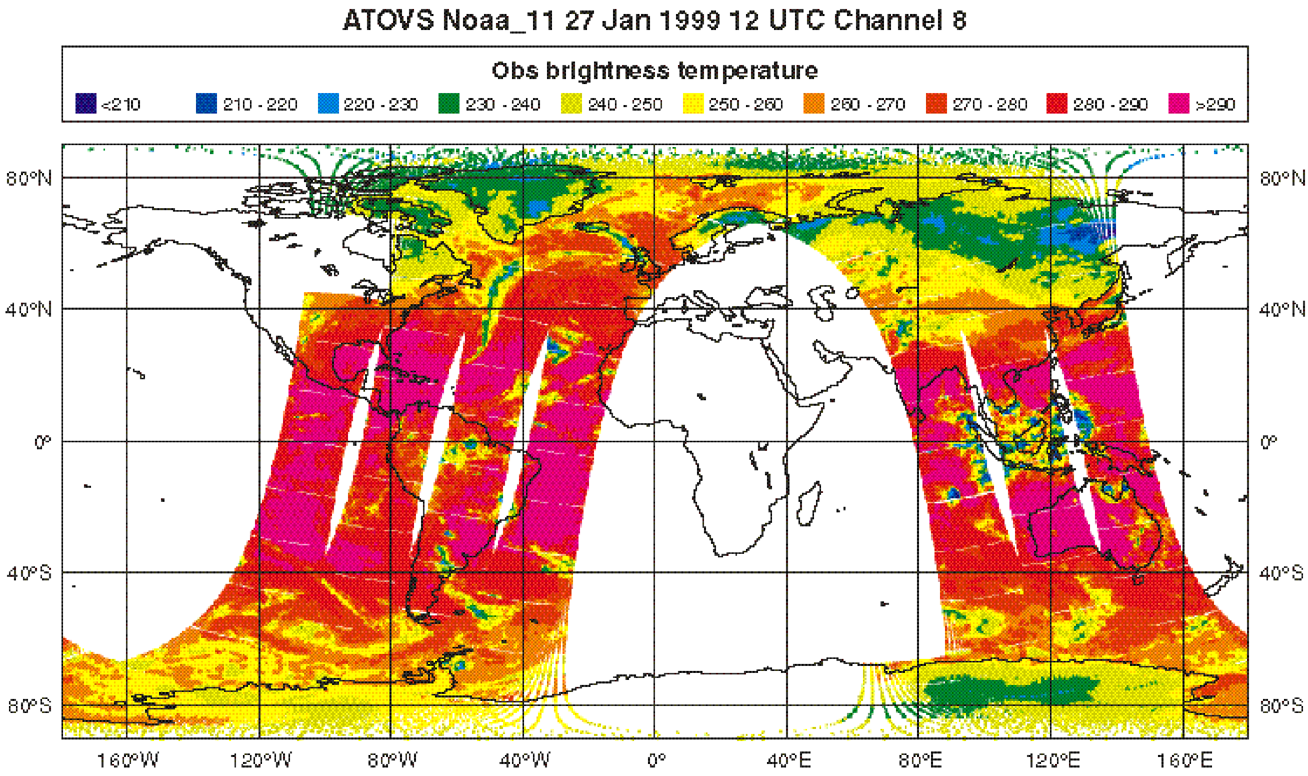

Figure 5: *Symbol plotting of TOVS Brightness Temperatures from orbiting satellite NOAA 11.*

longitude), by its position in centimetres from the bottom left hand corner of the subpage or by its X/Y position on a graph. There are facilities in symbol plotting that allow the user to control the height and colour of each symbol and, if required, a legend describing the symbols may be plotted.

Text facilities

MAGICS allows users to plot a block of text anywhere within the user's page. Atext block consists of a number of text lines, up to a maximum of ten lines, and may be positioned automatically by MAGICS or specifically by the user.

Facilities exist for plotting integer, real and character values, the necessary type conversion being done by MAGICS. It is also possible to plot the values of any MAGICS parameters in a text line.

Users have full control, via MAGICS parameters, over all aspects of text plotting, such as colour and style, as well as the height of text lines and the spacing between lines.

Instruction strings enable the plotting of complicated text strings and give even more user control of colour, height etc. Using Instruction strings, the user has control over each individual text character and may also plot symbols, e.g. the degree sign.

Legend plotting

MAGICS legend facilities enable users to annotate their plots. Legends can be produced automatically for contouring, wind plotting, satellite imaging, graphs and symbols. For contouring, attributes like style, colour and thickness are associated with a text which describes the relevant contour interval used in the plot. Shading legend entries consist of a sample of each shading pattern used in the plot. Contour legends can be one of two types, disjoint (default) or continuous. Examples of each type can be found in Figures 5 and 7 respectively.

For wind plotting, the shape and length of arrows are associated with a text which can describe the wind speed. Also, wind arrow colours may be associated with a text to describe the significance of the colours. Wind flags, plotted each time in the legend with a full barb, are associated with a text which describes the wind strength.

There are special legends for image plotting and in observation plotting of analysis feedback data.

The legend entries are plotted into an area known as the legend box, each legend entry consisting of three distinct parts: symbol, automatic text and user text. The user has full control over the positioning of the legend box, the number of legend entries and the way entries are plotted within the box.

There are examples of MAGICS legend facilities in Figures 2, 4, 5, 7 and on the front cover.

MAGICS shading

MAGICS shading allows users to shade areas in contouring, graph plotting, coastline drawing etc. with varying intensities, patterns and colours. There are parameters to control all the features and options in shading and their default values should ensure that, for most plots, a reasonable shaded plot can be achieved without having to set most of the shading parameters.

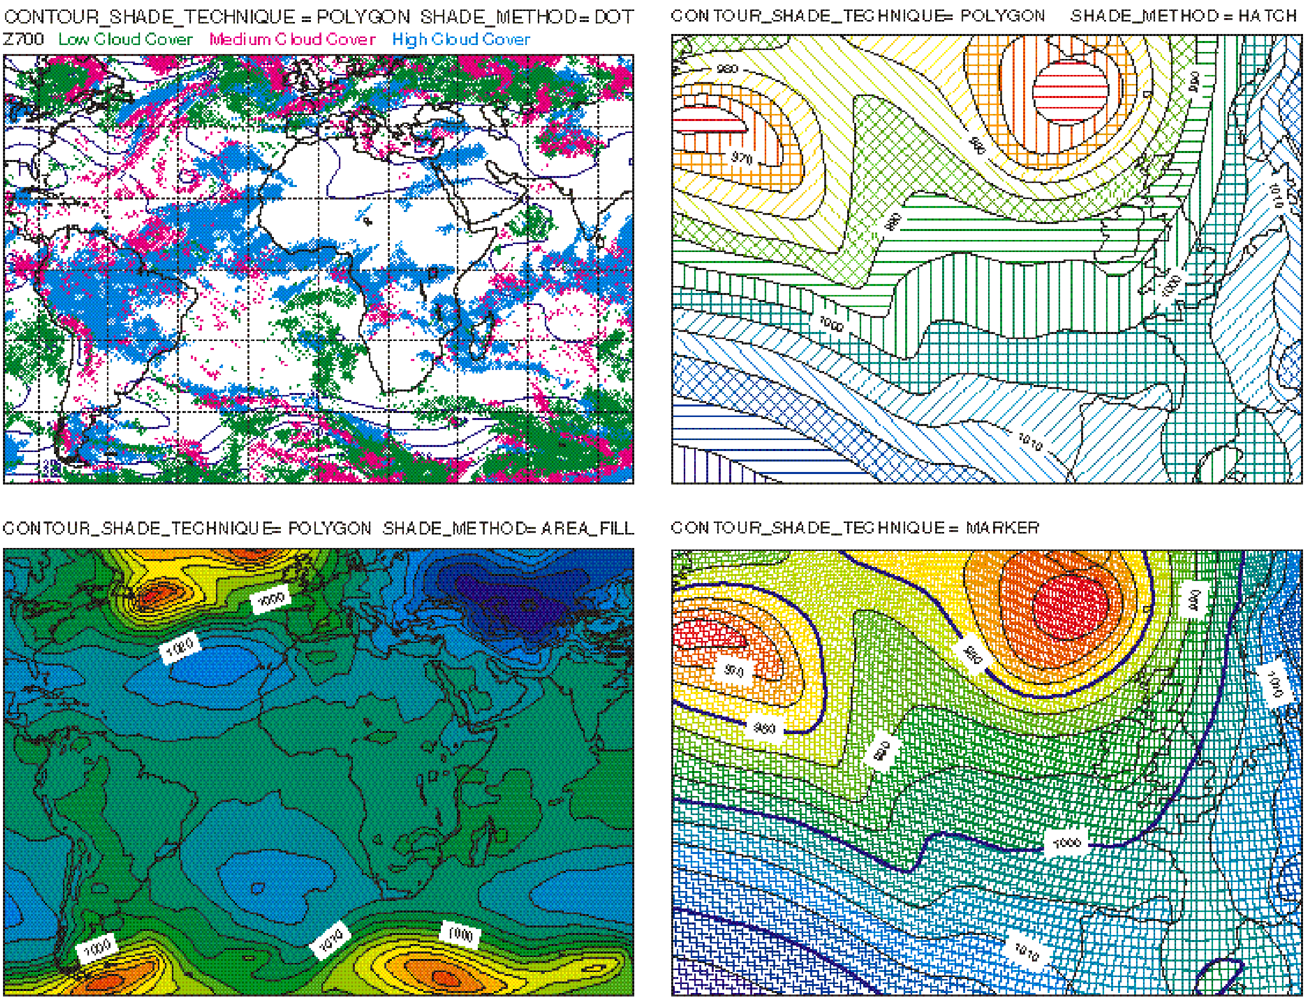

Figure 6: *Examples of MAGICS shading techniques and methods.*

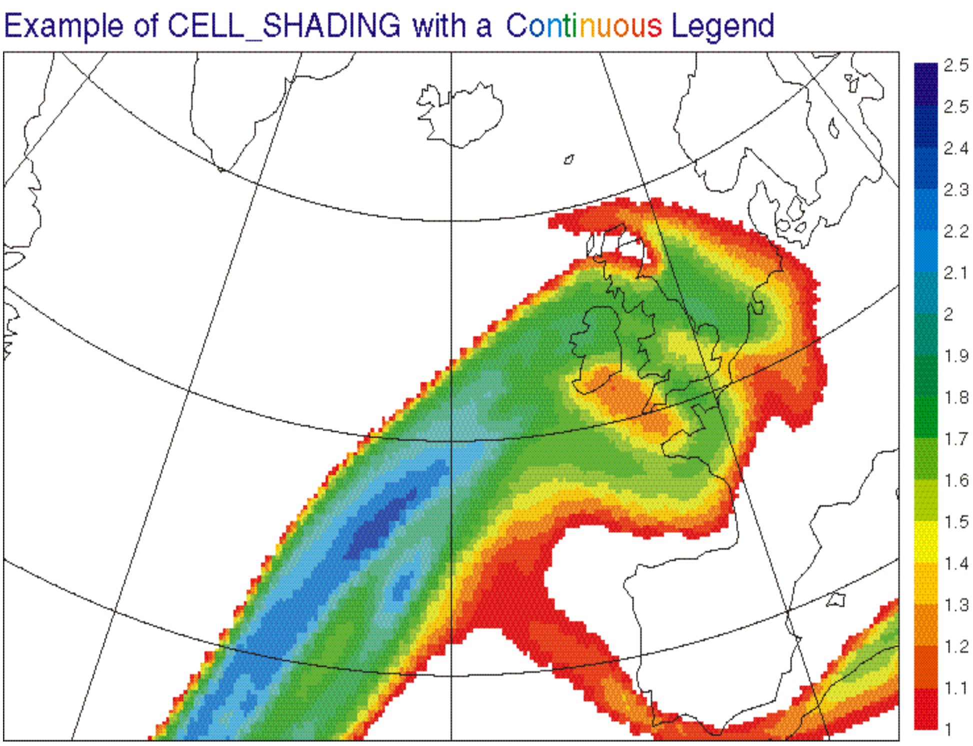

Figure 7: *Example using MAGICS CELL_SHADING technique and producing a CONTINUOUS legend.*

There are three different shading techniques used; polygon shading, tile (cell) shading and marker shading. Polygon shading is where the area between contours is formed into a closed polygon and filled. There are three different methods of shading available with polygon shading; dot, hatch and solid. Cell shading, which is a lower quality shading than polygon shading, shades each cell according to the value of the field at that cell. Marker shading consists of plotting a marker at each of the grid points in the field. The type of marker and also the height and colour of each marker is determined by the value at the grid point. Figures 6 and 7 give examples of the shading methods and techniques.

Shading enables more information to be plotted on a map without causing confusion. It is useful for highlighting specific details, e.g. negative temperatures, rainfall above 20 millimetres etc. Shading can make a map more easily understood, particularly when used with MAGICS legend facilities.

MAGICS Colours

MAGICS supports the use of colour for all graphical output, e.g. lines, text, shading. There are three ways to set a colour in MAGICS:

- The MAGICS colour name, e.g. RED;
- HSL (Hue, Saturation, Lightness) values;
- RGB (red, green, blue) values;

Specification groups

AMAGICS specification group consists of a set of MAGICS parameter values, where each group may consist of one or more parameter values.

Users may request MAGICS to save, retrieve or delete specification groups in memory. It is also possible to save and retrieve specification groups in a user file. If required, the user may alter this file by using normal text editors. MAGICS can read specification groups from a user file, which has been created by MAGICS or by the user.

Specification groups are useful when writing MAGICS application programs that produce, for example, a predetermined sequence of plots. Users can pre-define and save specific groups of MAGICS parameters that are used for each picture in the sequence. They simplify programming and make programs more readable.

Further development

MAGICS is a mature but still expanding graphics package. Future enhancements will be made in response to user's requirements and also to support the development of Metview.

Paddy O'Sullivan

ECMWF's computer status and plans

(This article is based on a talk given to all staff on 11 December 1998.)

Overview and status

The major changes over the past 12 months have been:

- termination of the CFS based data storage system and the removal of the IBM ES9000
- the installation of a 48 PE Fujitsu VPP700E system
- the replacement of all old PC systems with Windows NT based machines.

The configuration as of January 1999 is shown in Fig. 1.

The Fujitsu VPP700 has shown excellent availability throughout most of the year, with many weeks at 100% availability. Over the summer there was some disruption, partly due to hardware and software reconfiguring (in preparation for the 700E installation), and partly due to hardware failures. However, since late September availability has been very high again.

The utilisation of the Fujitsu VPP700 is now well over 80% (86% during the month of November for the non-I/O PEs). Work will be done during 1999 to see if some spare processor capacity on the I/O PEs can be utilised without impacting their I/O performance.

The utilisation of the Fujitsu 700E, installed 1 October 1998, has already exceeded 80%.

Revised Computer Division structure

Following the success in bringing the new Data Handling System into service, the project team has been disbanded and transferred into the regular section structure. At the same time the opportunity has been taken to rationalize some other aspects of the section structure. Thus the division now has 5 sections (previously 4), the new structure being:

Computer Operations
(Section Head: Claus Hilberg)
covering computer operation, operational support, and installation

Networking and Computer Security
(Section Head: Matteo Dell'Acqua)
covering LAN and WAN services, and computer security

Servers and Desktops
(Section Head: Richard Fisker)
covering servers, desktop systems and web services

Systems Software
(Section Head: Neil Storer)
covering Fujitsu and Data Handling services

User Support
(Section Head: Umberto Modigliani)
covering user advice, consultancy, training, documentation, libraries, liaison, etc.

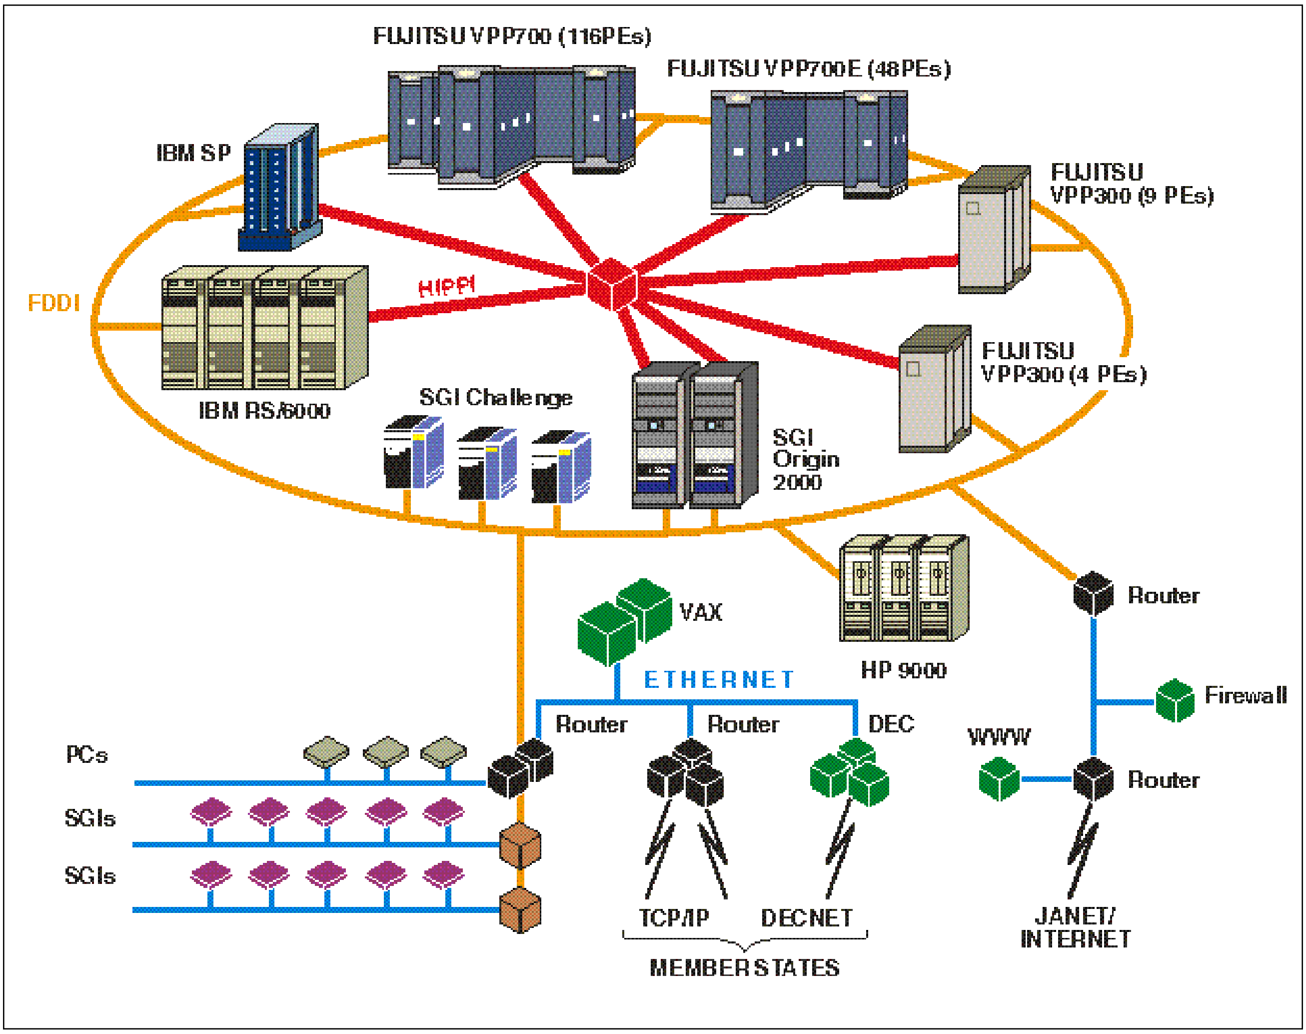

Figure 1: *The ECMWF computer configuration as of January 1999.*

Plans

High speed computing facility

The present contract with Fujitsu expires in December 2000. After intensive investigations of various options the Centre proposed to Council the contract be extended two years until December 2002. Council has now agreed to this proposal. Hence over the next three years the following will occur:

- December 1999 the capacity will be doubled by the installation of a new system. This machine will initially have about 40 processors (PEs), each approximately four times the power of a VPP700 series processor, and each with 4 Gbytes of memory.
- Mid 2000 the new machine will be increased to 100+ PEs. From this point, the aggregate performance of the Fujitsu systems installed at ECMWF will be about 400 Gflops sustained.
- It is planned to issue in early 2001 an Invitation To Tender (ITT) for a replacement.
- December 2002 the contract ends.

Throughout this period the present 116 PE VPP700 and 48 PE VPP700E will remain in service.

Workstations

The existing purchase agreement with Silicon Graphics terminates in mid-1999. Before embarking on an ITT for a replacement the Centre will thoroughly review its desktop strategy. As well as taking user requirements, especially those for visualisation, in account, alternatives to the present Unix mainstream systems will be investigated, e.g. Windows NT, Linux, X-Terminals, Java stations, etc.

The ITT itself will be issued late in 1999 with the first of the new systems being installed in mid 2000.

Data Handling

A lot of effort went into migrating data from CFS to the new system, CFS was then finally terminated on 21 September 1998.

The new system is working quite well, but is taking a substantial amount of human resources in its day-to-day management. This issue will be addressed during 1999. In addition, extra hardware (tape drives, robots, processor capacity) will be installed to cope with the estimated increase in load.

A disaster recovery system has been ordered and will be brought into operation during 1999. This will hold copies of crucial data and files at a separate location and on a different media technology from the present DHS.

Local Area Network (LAN)

Quite a lot of the present LAN equipment is either out of date or overloaded and would certainly not cope with the increased load from the new Fujitsu system to be installed in late 1999. Hence an ITT for a refurbishment of the local area network will be issued early in 1999.

Recently LAN problems have been seen, especially for users outside the main office block at ECMWF. It is hoped to be able to resolve these problems soon, without waiting for the outcome of the ITT.

To help users optimize their work when it comes to moving large amounts of data around, a series of guidelines on the best ways to use the network utilities (e.g. rcp, NFS, etc.) will be issued.

Wide Area Network (WAN)

Acontract was recently signed with EQUANT to provide the so-called RMDCN network. This network will:

- replace all the existing lines between ECMWF and its Member States
- provide GTS connections between ECMWF's Member States
- provide GTS connections for up to 49 countries within WMO's Region VI.

Apilot project will start early in 1999 with acceptance of the new network scheduled for July 1999.

Miscellaneous

Among the other planned activities are:

- an upgrade of the HP servers to cope both with an increasing amount of incoming satellite data and the transfer of product dissemination to the HP servers from the VAX systems
- Novell Netware will be phased out, to be replaced by Samba
- a web based PREPIFS service for Member States
- the setting up of a Year 2000 testbed in the second quarter of 1999. This is to allow all crucial software to be tested on a system running with the dates past 31.12.1999.

Andrew Lea/Walter Zwieflhofer

CFS terminated - the end of an era

One of the central roles of ECMWF is to "collect and store appropriate meteorological data".

When ECMWF first started to store such 'appropriate' data (late 1970s) the only suitable medium was 12" open reel magnetic tape. By the early 1980s 4000 such tapes were being added each year, using manual systems both to handle them and keep track of their contents, and the demand was increasing. Clearly these manual systems could not cope much longer, so ECMWF looked for a major data repository system that would automatically handle large amounts of data. In 1983 ECMWF started to work with the CFS file system. CFS, or the Common File System, had been developed by the Los Alamos National Laboratories (LANL) to handle the large amounts of data it was required to store.

Using CFS an operational service was started at ECMWF in 1984. It provided ECMWF with all its data storage needs for almost 15 years, yet it too eventually became unsustainable. Hence at 10.01 UTC on 21 September 1998 the CFS service was finally terminated.

This article is an attempt to review that service, and the central role CFS has played in maintaining ECMWF's data storage service for all those years.

In the beginning . . .

ECMWF's initial file handling system consisted solely of manually mounted open-reel tapes. This system first became operational in 1978. By 1983 the number of tapes in the manual library had grown to around 20,000, and 4000 more were being added each year. Further expansion of that system was clearly insupportable, and a major research and evaluation exercise was launched to find an automated alternative.

As a result, the Centre installed in 1983 the (then) Common File System of Los Alamos National Laboratory (Collins[1]; Mclarty[2]), running on a dedicated IBM processor.

CFS was a hierarchical data storage system that attempted to optimize retrieval times by storing active files on fast access storage devices (disks) and large, less active, files on slower, less expensive devices (cartridges). Infrequently used files were stored off-line on magnetic tape. Files were automatically moved between the various storage media by a file migration program which analyzed file activity and size, matching it to storage device capabilities.

The Common File System (CFS) had developed in a unique environment, however, and a number of additions and modifications were necessary before it could successfully be deployed in the Centre.

ECMWF and Los Alamos staff therefore worked together on the system for over a year, adding features and providing a different network interface. The system entered production at the end of 1984.

The hardware chosen to run the system initially was an IBM 4341 processor with 12.5 Gbytes of on-line disk space, 35.5 Gbytes of on-line mass storage (IBM 3850 cartridges), and off-line 6250 bpi tapes.

From the start, a major objective was to provide a meteorology-oriented view of the stored data. CFS in its original form offered a 'bottomless pit', cost-free file service to the

Los Alamos user community. We knew that, given total freedom from constraints, anarchy would rapidly develop in the Centre's user community, leading to the establishment of private data archives and much duplication of data.

Therefore, in parallel with the developments in CFS itself, a major project was established to provide an organised, centrally managed, application-oriented interface to the primary archives. This project was called MARS (for Meteorological Archive and Retrieval System; Hennessy[3]). MARS provides a meteorologist's interface to the stored data, delivering sections and concatenations of stored data objects as required in a form suitable for meteorological research. The success of the MARS project has provided a major benefit to all of the users of ECMWF's services.

As CFS became more popular outside its original environment, an attempt was made to market it commercially as the 'Datatree' product which was promoted by General Atomics 'Discos' division. No commercial sales were ever made (probably because of the need to establish a full-scale mvs installation to support the product), and Discos went on to promote and market the rival Unitree system instead.

Enhancements to the original CFS

We determined that, for our purposes, changes were needed to CFS in a number of areas. Six major areas were addressed to prepare it for service in the Centre's environment. In all of these areas we received considerable help and support from Los Alamos staff, to the extent that they provided the majority of the effort in some cases.

The main areas addressed in this effort were:

Network Interface

- The original CFS used a home-grown network based on SEL minicomputers. We needed to connect it, and its client interfaces, to a commercially-available network.

Magnetic tape support

- We needed support for conventional half-inch open-reel tapes, for reasons of interchangeability and archival permanence.

Magnetic tape volume families

- The concept of tape volume 'families'was added to permit better localisation of reference in the managed archives.

Partial get

- A facility to request and deliver only part (or possibly many parts) of a larger dataset was imperative for the applications we had in mind.

New client packages

- At that time, CFS had no client package available to run on the server platform, or on two of our critical systems.

The network interface

The choice of network medium and software was not a simple one in 1982; the more so since we had no UNIX systems on our site. We selected the CDC LCN/RHF network, a 50 Mb/s CSMA/CD LAN. Unfortunately, the CDC system had no application interface (API) available on the IBM platform which would run CFS, so our first task was to design and add the missing API and cross-memory services needed to move data from the network manager to the CFS subsystem. Further functional and conceptual difficulties arose which necessitated protocol changes.

The changes to CFS source code were kept as localised and formalised as possible through this process. In fact, LANL helped us greatly, by accepting the majority of our changes – in this area and others – back into the source libraries which they maintained.

Although the result was in the end functionally satisfactory, we were only able to obtain data rates approximately 20% of those seen by LANL with a similar class of processor.

Magnetic tape support

Our Centre needed the archival permanence and data exchange capability of a standard and widely accepted data storage medium. In 1982 there was none available in the industry, apart from the conventional half inch open-reel magnetic tapes which were already in use in the manual system. CFS in its native habitat had been based on a storage hierarchy of three layers: real disk; on-line virtual disk provided by an IBM 3850 Mass Storage System; and off-line virtual disk consisting of racked cartridges for the 3850.

Fortunately, it had become clear at the time that magnetic tape support would be needed by LANL themselves, and the work to add it was carried out jointly between ourselves and LANL. At about this time, the IBM 3480 cartridge tape drive became available. LANL determined that they would use the new medium immediately, whereas we remained with conventional tapes for some years. Later, however, our system used cartridge tapes exclusively, mounted in four StorageTek tape silos.

Magnetic tape volume families

A primary objective of the MARS system design was, as far as was possible, to capitalize on the expected locality of reference to data contained in off-line tapes. We had seen from analyses of the old system that, once a tape had been mounted, the probability of its being accessed again within the measurement period rose immediately, and continued to rise with every successive access. We needed to preserve - and if possible to improve upon - the locality of related data which gave rise to this characteristic behaviour.

The simple scheme adopted was to define 'families'of files whose access patterns were expected to be correlated. Families were uniquely named, and the family designation attached to the directory tree at some sub-directory (non-leaf) node. All leaf nodes subsequently added to the directory tree below this level (corresponding to files being created) were then members of the family. Files which were members of a family could be placed in an empty tape volume or on one which was logically affiliated to the family; and writing a family file to an empty tape volume immediately affiliated it to that family. The effect therefore was that one could isolate a group of tapes which contained only files from a particular directory sub-tree - and were therefore, presumably related in some way. A further requirement, to obtain useful locality, was that files created contemporaneously should occupy adjacent locations on the tape. We did not put in place any formal systems to achieve this, since the normal action of the archive-to-tape subsystems tended to have this effect in any case. However, in some circumstances this relationship could break down and cause problems.

Partial get

The data objects in the system spanned a wide range of sizes - from 6 kilobytes to 250 kilobytes. The top of this range was still rather small compared with the file sizes which could be economically stored on the available mass storage devices. So long as the system retained an equivalence between logical and physical files, a mechanism was clearly needed to efficiently concatenate data objects into larger files, and later extract parts (or more generally, a collection of parts) of the file for delivery to an application. To make this process efficient demanded some restrictions on the alignments and sizes of data objects and extraction sets.

We therefore put in place a simple implementation of partial access on byte boundaries, which accepted, along with the request, a vector of skip/copy count values. Unfortunately, delivering the extraction vector to the archive system and associating it with the correct request turned out to be a non-trivial matter. Several different schemes were adopted, with varying degrees of success.

New client packages

The only CFS client implementation available to us when we started was for our CDC mainframes running NOS/BE. Architectural features of that package made it entirely unsuitable for re-implementation on any other platform, so that our initial efforts were dependent on using the CDC machines as gateways from application systems onto the file system. This was clearly unsatisfactory in anything but the short term.

However, during the time we were installing CFS, LANL was working on a new request protocol. We took the new protocol implementation (which was still experimental), and constructed a client implementation to work with it. This was intended to be as portable as possible, so that we would be able to have client implementations on all our critical target machines. The first platform we chose was VAX VMS, after which we ported the client to IBM MVS: for the first time, applications running on the file server system itself had access to CFS files.

Later we constructed a TCP/IP-based interface using standard Telnet and FTP. An X-based front-end was also added. Thus finally we had access to CFS from all of our general purpose machines.

And towards the end . . .

Although originally supporting only the MARS service, support for a user's own private files was added in 1987. It was known as ECFILE. Access to ECFILE from most of the other systems at ECMWF was via a special command that allowed a subset of the native CFS commands to be used.

Over the years the hardware on which CFS ran was enhanced and upgraded several times. The IBM 4341 was replaced by an IBM 3090, and the 3850 tape cartridges plus the 6250 bpi off-line tapes by IBM's new 3480 cartridge tape system. The 3480 cartridges initially had a capacity of 200 Mbytes, later through upgrades and then hardware compression, this was raised to between 600 and 800 Mbytes per cartridge. StorageTek 4400 silos with 4490 tape units were added to handle the ever growing cartridge library. At the end there were four such silos attached.

ECMWF's mass-storage system based on CFS was very successful. Given that CFS was designed originally to use LANL's home-grown network hardware and protocols (this was long before TCP/IP), and that it depended exclusively on the IBM 'steam driven'3850 robotic mass storage unit with 50 MB cartridges, its ability to adapt to subsequent technology and requirements was remarkable. The original design of CFS was intended to handle of the order of one terabyte of data and one million files. At the end, ECMWF had CFS worked reliably in conjunction with TCP/IP over HIPPI channels, handling over 65 terabytes of data, 12 million files, 85,000 tape volumes, and traffic in excess of 100 GB per day. Availability of the system was of the order of 99%.

However, in the present environment at ECMWF, CFS finally showed its age. Although the IBM mainframe and its MVS operating system gave an extremely reliable service, it was an uneasy fit in an environment where the client machines are almost universally running UNIX variants of one sort or another, and in which all the clients were using TCP/IP. For example, there was effectively no possibility of providing a mounted-file interface (NFS; AFS, DFS) to CFS. Furthermore, the structure of CFS - although highly reliable - did not make it easy to adapt the system to new storage hardware. And, finally, LANL itself decided that its future data storage needs could not be met with CFS or any development of it.

Conclusions

The CFS based mass data storage system, and its application-specific interface MARS, was an operational success and contributed greatly to the Centre's operations. However, this success was bought at a considerable price in manpower investment to provide local enhancements and application-specific features in the storage server.

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

As said before one of the prime reasons for the success of CFS at ECMWF has been the excellent partnership between LANL and ECMWF. Among the several LANL staff involved thanks must go to one individual, Emily Willbanks, who worked extensively with us in providing very good support in the joint development required to make CFS suitable for our needs.

The success of CFS was also based on the reliability of the software and hardware on which it was running, namely the IBM MVS operating system, the IBM hardware, and the StorageTek tape silos.

Let's hope the new Data Handling System will enjoy the same success in meeting the challenge that ECMWF faces with the ever growing demands for data storage over the next 15 years.

References

1. **Collins**, **B.**, **Devaney**, **M.**, and **Willbanks**, **E.** (1982). Anetwork file storage system. *Digest of Papers, Fifth IEEE Symposium on Mass Storage Systems*, Boulder, Colorado, October 1982.
2. **McClarty**, **T.**, **Collins**, **B.** and **Devaney**, **M.** (1984). A functional view of the Los Alamos central file system. *Digest of Papers, Sixth IEEE Symposium on Mass Storage Systems*, Vail, Colorado, June 1984.
3. **Hennessy**, **J. T.** (1986). MARS - The ECMWF Meteorological Archive and Retrieval System. *Proceedings, 66th annual meeting of the American Meteorological Society*, Miami, Florida, January 1986.

Dick Dixon

ECMWF software and the year 2000

In common with many organisations, ECMWF is assessing the impact the Year 2000 will have on its software. In particular, a Year 2000 Task Group has been set up to identify, study, and address the problems which are likely to occur to its computer hardware and software at the 1999/2000 and 2000/2001 boundaries, and at the leap year date of 28/29 February 2000.

Information concerning the progress of the Year 2000 work, the status of the software available to Member States, and the availability of test data sets has now been collected together on web pages accessible by Member States. These pages can be found at:

http://www.ecmwf.int/ecmwf/info/y2k/

They will be updated as required.

The following items should be noted in particular:

ECLIB date routines

New date handling routines have been produced in ECLIB. The old routines are scheduled to be removed from the library on 10 February 1999 as reported in ECMWF Computer News Sheets 352, 354 and 355.

Software releases

For all software which is not obsolete, ECMWF will produce a version adapted for ECMWF's use in the Year 2000. Alist of the current status of ECMWF software in respect of Year 2000 will be maintained on the ECMWF web server. Software marked 'ready for testing' can be made available.

Test data

ECMWF has produced some data sets for testing its own operational system with respect to the Year 2000 problem. These will be made available to Member States, and their location on the ECMWF computers will be given in the Year 2000 web pages. These data sets will be added to later.

Please contact User Support in the first instance if you have questions about any aspect of this project. If you require updated software, please contact ECMWF's Data Services.

Umberto Modigliani

ECMWF annual seminar

Diagnosis of Models and Data Assimilation Systems, 6-10 September 1999

The practical importance of weather forecasts, and concerns about environmental change, have made the development and diagnosis of atmospheric and ocean models, and their associated data assimilation systems, a key focus for world-wide meteorological and oceanographic science. There is close similarity between the models used for work on medium and extended range forecasting and the general circulation models (GCMs) used for climate research. Both modelling communities rely heavily for model validation on the extensive climate datasets prepared by national and international agencies under the aegis of the WMO/ICSU. In addition, model errors can be diagnosed when they grow in an otherwise accurate series of forecasts that start from accurate analyses.

Further diagnostics stems from the data assimilation systems needed for forecasting practice and from verifications of operational forecasts and analyses. The seminar will provide a pedagogical review of current methods to diagnose and resolve problems in models and assimilation systems, to identify missing processes in the models and to refine the representations of existing processes.

Posters providing further information on the programme and application forms will be distributed around May 1999.

Els Kooij-Connally

WCRP

2ND International Conference on Reanalyses

CALL FOR PAPERS

World Climate Research Programme • European Centre for Medium-Range Weather Forecasts
National Aeronautics and Space Administration • National Oceanic and Atmospheric Administration
National Science Foundation • American Meteorological Society
International Business Machines • Fujitsu Systems (Europe) Limited

Second International Conference on Reanalyses

23 to 27 August 1999

Wokefield Park, Mortimer, Reading, UK

Reanalyses of the instrumental atmosphere ocean and land data record, produced by modern data assimilation methods, provide a key resource for many activities of the World Climate Research Programme, in particular the atmospheric modelling studies of the Working Group on Numerical Experimentation (WGNE), and the investigations being undertaken within the Global Energy and Water Cycle Experiment (GEWEX), the World Ocean Circulation Experiment (WOCE) and the Climate Variability and Predictability (CLIVAR) programme. Fundamental topics of interest include:

– the deepening of our understanding of low frequency variability

– the development and validation of coupled atmosphere – ocean – land – ice – chemistry models

– the assessment of predictability, spanning the medium-range time-scale through seasonal and inter-annual time-scales to decadal time-scales

The Second International Conference on Reanalyses will assess progress and problems in exploiting reanalyses for these purposes. It will also identify the main priorities for further work in exploitation of existing reanalyses and for the preparation of new reanalyses.

The subject areas to be considered will include:

- Global reanalysis datasets, Analysis methodology, Impact of data inhomogeneities, Estimation of data bias, Evaluation against independent observation data sets
- Diagnosis from reanalyses of Global and Regional Aspects of the Hydrological Cycle, the Energy Cycle, and the Angular Momentum Cycle
- Diagnosis from reanalyses of the mechanisms of atmosphere and ocean phenomena such as the Madden-Julian oscillation, stratospheric variability, inter-annual variability in Monsoons and tropical cyclone frequency, and ENSO variability
- Use of Reanalyses for driving ocean, land, sea-ice and chemistry models, and for development and validation of Climate System Models and their sub-systems
- Medium-range, Seasonal-to-Interannual and Decadal Predictability Studies using Reanalyses
- Future Requirements and Plans

The format will include two invited talks per day, with normal presentations lasting 15-20 minutes. The poster sessions will be social occasions in congenial surroundings.

Wokefield Park is a modern fully-equipped conference centre set in parkland a few miles from Reading and featuring an eighteen hole golf course.

See http://www.ecmwf.int for further information. Submission of abstracts is by electronic means to this site.

For specific queries contact the chair of the Local Organising Committee:

Dr. D. Burridge
ECMWF, Shinfield Park, Reading RG2 9AX, UK
Tel: +44 118 949 9000 Fax: +44 118 986 9450
e-mail: Reanalyses2@ecmwf.int

Member State computer resource allocations 1999

Member State	Fujitsu (kunits)	Data (Gbytes)
Belgium	126	801
Denmark	103	651
Germany	614	3895
Spain	206	1303
France	424	2684
Greece	91	576
Ireland	77	489
Italy	359	2272
Yugoslavia*	82	523
Netherlands	155	979
Norway	98	621
Austria	117	743
Portugal	90	568
Switzerland	135	858
Finland	90	573
Sweden	121	768
Turkey	109	690
United Kingdom	342	2170
Special projects	371	1836
Total	3710	23000

* In accordance with UN Security Council Resolution 757 (1992) of 30 May 1992, the Council instructed the Director to suspend the telecommunications connection to Belgrade with immediate effect. This took place on 5 June 1992. As a consequence no operational products are disseminated to Belgrade and access to the Centre's computer system is not available to Belgrade.

Special Project allocations 1999

Member State		Institution	Project title	Fujitsu units	Data storage Gbytes
Continuation Projects					
Austria	**1**	Univ für Bodenkultur, Vienna (Kromp-Kolb)	Vertical ozone transports in the Alps	500	4.2
	2	Univ Vienna (Ehrendorfer)	Covariance evolution and forecast skill	930	6.0
	3	Univ Innsbruck (Mayr)	Heavy convective precipitation over and along mountains – numerical simulations	1,700	3.0
Belgium	**4**	Univ Louvain (Van Ypersele)	Modelling the climate and its evolution at the global and regional scales (CLIMOD Network)	20,000	500.0
France	**5**	L.A.M.P., Aubière (Cautenet)	Chemistry, cloud and radiation interactions in a meteorological model	93	2.0
	6	CERFACS (Thual)	Universal software for data assimilation: variational method for global ocean	1,680	15.0
	7	CNES (Courtier)	MERCATOR	58,272	220.0
Germany	**8**	Inst. for Geo. and Met. (Speth)	Interpretation and calculation of energy budgets	93	6.0
	9	MPI, Hamburg (Bengtsson)	Numerical experimentation with a coupled ocean/atmosphere model	18,000	85.0
	10	MPI, Hamburg (Bengtsson)	Simulation and validation of the hydrological cycle	15,000	75.0
	11	Univ. of Munich (Wirth/Egger)	The behaviour of cut-off cyclones in ECMWF analysis: impact of diabatic processes on their development and decay	70	1.0
	12	FU, Berlin (Fischer)	Comparison of the ECMWF cloud scheme with multi-spectral satellite data in the Baltic area	1,500	7.0
	13	GKSS, Geesthacht (Rockel)	Energy and water cycle components in regional forecasts, remote sensing and field experiments	10	0.2
	14	Univ Munich (Stohl)	Validation of trajectory calculations	200	3.0
	15	DLR, Wessling (Hoinka)	Climatology of the global tropopause	50	10.0
Ireland	**16**	Met Éireann (Lynch)	The HIRLAM 4 project	3,500	5.0
Italy	**17**	Ist. per lo Studio della Dinamica delle Grandi Masse, Venezia (Cavaleri)	Testing and applications of a third generation wave model in the Mediterranean Sea	3,000	3.0
	18	Univ Bologna (Rizzi)	TOVS 1B radiances and model simulations	300	10.0

Member State		Institution	Project title	Fujitsu units	Data storage Gbytes
Netherlands	**19**	KNMI (Siegmund)	Stratosphere-troposphere exchange	2,000	8.0
	20	KNMI (van Velthoven)	Chemistry and transport studies with a 3D off-line tracer model	3,000	25.0
	21	KNMI (Drijfhout)	Agulhas	15,000	0.0
	22	KNMI (Komen)	Validation of re-analysed A/S fluxes	20,000	5.0
	23	KNMI (Siebesma)	Large Eddy Simulation (LES) of boundary layer clouds	10,000	10.0
	24	KNMI (Opsteegh)	Short term, regional probabilistic forecasting using IFS	5,000	12.0
	25	KNMI (Kelder)	The relation between satellite ozone observation errors and dynamical errors in the ECMWF model	1,700	10.0
	26	KNMI (Burgers)	OGCM mixed-layer modules and assimilation	15,000	100.0
Norway	**27**	Geoph. Inst., Univ. of Bergen (Grønås/Kvamstø)	Parametrization of clouds in general circulation models	70,272	0.2
	28	Univ Oslo (Isaksen)	Ozone as a climate gas	1,000	5.0
Switzerland	**29**	Univ. Berne (Schüpbach)	The climatology of global transport (long-range) to 'background'stations	10,000	85.0
UK	**30**	Univ Reading (Hoskins)	Routine back trajectories	3,000	4.0
	31	Br Antarctic Survey, Cambridge (Turner & Leonard)	Assessment of ECMWF forecasts over the high latitude areas of the S. Hemisphere	0	1.0
New projects					
Austria	**32**	Univ Vienna (Dorninger)	Estimating the energy budget climatology over the Alps from different data sources	30	5.0
France	**33**	CERFACS (Terray)	Decoded variability over the North Atlantic – European region	10,000	150.0
	34	CERFACS (Rogel)	Seasonal to interannual predictability of a coupled ocean-atmosphere model	10,000	150.0
Germany	**35**	MPI Hamburg (Manzini)	Middle atmosphere modelling	13,000	50.0
Italy	**36**	CINECA, Bologna (Molteni)	Non-linear aspects of the systematic error of the ECMWF coupled model	20,000	30.0
			Total allocated	**333,900**	**1,605.6**
			Reserve (to be allocated by ECMWF)	**37,100**	**230.0**
			TOTAL	**371,000**	**1,835.6**

Index of past newsletter articles

This is an index of the major articles published in the ECMWF Newsletter series. As one goes back in time, some points in these articles may have been superseded. When in doubt, contact the author or User Support.

	No	Date	Page
GENERAL			
ECMWF publications – range of	74	Winter 96/7	21
COMPUTING			
Call Desk	71	Winter 95/6	16
Data handling system			
– *IBM RS6000 based*	70	Summer 95	15
– *New Service*	78	Winter 97/8	8
ECMWF Documentation			
– *Current Computer Bulletins*	80	Summer 98	22
(Exabyte)			
8mm data cartridge service	67	Autumn 94	36
Fortran 95	73	Autumn 96	31
Fujitsu VPP700	74	Winter 96/7	17
&	76	Summer 97	14
High performance Fortran	78	Winter 97/8	8
MAGICS – the ECMWF graphics package	71	Winter 95/96	8
MARS			
– *the ECMWF archive/ retrieval system*	72	Spring/ Summer 96	15
– *efficient use of*	72	Spr/Sum 96	21
– *implementing*	75	Spring 97	9
METVIEW	68	Winter 94/5	9
Multitasking ECMWF spectral model	60	Dec 92	3
Networks			
– *ECMWF's internal network*	54	Jun 91	20
– *New LANs at ECMWF*	59	Sep 92	20
– *ECNET*	71	Winter 95/6	15